arXiv:1501.00324v1 [cs.MS] 1 Jan 2015

A New Sparse Matrix Vector Multiplication GPU Algorithm Designed for Finite Element Problems

J. Wong,[1] E. Kuhl,[1] E. Darve[1,2]

[1]Mechanical Engineering Department,
[2]Institute for Computational and Mathematical Engineering,
Stanford University

January 5, 2015

SUMMARY

Recently, graphics processors (GPUs) have been increasingly leveraged in a variety of scientific computing applications. However, architectural differences between CPUs and GPUs necessitate the development of algorithms that take advantage of GPU hardware. As sparse matrix vector multiplication (SPMV) operations are commonly used in finite element analysis, a new SPMV algorithm and several variations are developed for unstructured finite element meshes on GPUs. The effective bandwidth of current GPU algorithms and the newly proposed algorithms are measured and analyzed for 15 sparse matrices of varying sizes and varying sparsity structures. The effects of optimization and differences between the new GPU algorithm and its variants are then subsequently studied. Lastly, both new and current SPMV GPU algorithms are utilized in the GPU CG Solver in GPU finite element simulations of the heart. These results are then compared against parallel PETSc finite element implementation results. The effective bandwidth tests indicate that the new algorithms compare very favorably with current algorithms for a wide variety of sparse matrices and can yield very notable benefits. GPU finite element simulation results demonstrate the benefit of using GPUs for finite element analysis, and also show that the proposed algorithms can yield speedup factors up to 12-fold for real finite element applications.

1. INTRODUCTION

The recent use of graphics processors (GPUs) for scientific applications has opened up possibilities in achieving real-time simulations, and has been shown to provide remarkable speedup factors for many different applications. Scientific applications using GPUs range from computing the flow over hypersonic vehicles [1] to finite element simulations of virtual hearts [2] and of viscoelastic properties of soft tissue [3] to medical applications [4, 5, 6]. There is a developing body of literature concerning GPU implementation of finite element analysis codes.

Early papers investigating the use of GPUs for scientific computing focused on conjugate-gradient (CG) and multigrid solvers [7, 8, 9]. Several studies in the area of finite-element analysis (FEM) have focused on the solution of the sparse linear system of equations $Ax = b$ resulting from a FEM discretization [10, 11, 12], mainly because the solution stage is often the most computationally intensive step. Some assembly strategies for the GPU have been mentioned as well. Often, specific applications have allowed special approaches for FEM assembly. The methods described in [12] for geometric flow on an unstructured mesh and in [13] for FEM cloth simulation derive relatively simple expressions for each non-zero in the system of equations. The relative simplicity and the inherent parallelism of computing each non-zero independently makes these approaches well suited for GPUs. The PhD thesis of Göddeke [14] has an extensive discussion of the history behind GPU computing and the application of multigrid to FEM and its optimization for GPUs. Applications to the FEM package FEAST [15] are discussed in [14].

Formulations have been considered in connection with discontinuous Galerkin methods. See [16] where the method is applied to Maxwell's equations. For high-order, continuous Galerkin methods, an assembly strategy is proposed in [17, 18]. Other successful approaches, such as [19] for deformable body simulation, took advantage of vertex and fragment shaders on the GPU to perform the assembly. New hardware and a more general computing environment appear to have made some of these techniques obsolete. Markall et al. [20] compare FEM implementations on CPUs and GPUs. This paper includes a discussion of FEM along with spectral element methods and low-order discontinuous Galerkin methods. Recent investigations include [21, 22, 23, 24].

The difficulty in optimizing FEM calculations for GPUs has led some groups to develop special purpose languages and compilers to automate part of the process and separate the task of implementing novel numerical methods from the task of optimizing the computer implementation. [25] discusses the Unified Form Language (UFL) for this purpose.

While GPUs are already used for many inherently parallel operations, such as material point O.D.E. solvers for finite elements [26, 27] and for finite differences [28], total performance gains for finite elements on GPUs can only be realized by efficiently combining all aspects of the calculation: element calculations, global assembly, and solver.

In this work, one of focuses is the sparse matrix solver that is used as part of FEM. However, it is generally not possible to simply "port" CPU algorithms to the GPU. Typically such implementations result in a GPU code that may be slower than the CPU code. Differences in GPU memory architecture, speed of individual GPU cores, and the communication between GPU cores, GPU memory, and host (CPU) memory can greatly impact the performance of GPU algorithms. Because of these differences, much effort has been spent in developing new algorithms designed for general purpose use.

The performance of many of these algorithms is dependent on the structure of the data. In particular, sparse matrices can be represented by a variety of formats and the computational performance may differ greatly depending on the format used [29]. There is a large body of literature on the investigation of various formats and their efficiency.

A sparse matrix contains mostly null entries. As a result it is advantageous to store only the non-zero entries along with their location in the matrix. The difficulty with such a storage is that it leads to a random access to memory, such that the performance is heavily limited by memory bandwidth and often very low. As a result several storage schemes attempt to make the memory access much more regular such that the effective memory bandwidth increases. This is crucial on GPUs where coalesced [30] memory access (i.e., accessing a 128-byte aligned segment in the device memory) is crucial for best performance. Bell et al. [31, 29] was an early paper on SMPV and discusses all the key sparse matrix formats including the diagonal format (DIA); ELLPACK, in which all non-zero entries are stored in an $N \times K$ dense matrix (N is the number of rows in the matrix and K the maximum number of non-zeros per row); the coordinate format (COO, with a simple `row`, `col`, `data` format); compressed sparse row format (CSR, this is a popular format on CPUs); and the hybrid format that combines ELLPACK with COO. At the end of the paper, the packet format (PKT) is discussed. In this format the matrix is decomposed into blocks with a high density of non-zero entries and a type of "local" CSR format is used to store entries associated with each block (offsets from the base index are stored using 16-bit integers).

Before exploring in more details various papers that followed the footsteps of Bell, we mention the CUSP library [32] by Bell and Dalton which implements various matrix formats and iterative schemes. cuSPARSE [33] is a library released by NVIDIA which contains code for SPMV, SPMM (sparse matrix-matrix addition and multiplication), sparse triangular solve, a tri-diagonal solver, and incomplete factorization preconditioners.

Among the newer implementations, we start with a series of papers on variants of the ELLPACK format. ELLPACK-R [34] uses an ELLPACK format and each thread computes a single row. An additional array is used to store information about the true number of non-zero elements per row (not accounting for the ELLPACK padding), such that each thread is not required to do more floating-point operations than strictly necessary. This reduces the number of extra flops usually found in ELLPACK. Vazquez [35] revised his original algorithm and renamed it ELLR-T in 2010. In this work, he uses T threads to compute a single row in SPMV. The Sliced ELLPACK [36] format is similar to ELLPACK. However it is more flexibile and uses a storage scheme that varies the number of non-zero entries stored across CUDA warps. This amounts to an ELLPACK where K is a function of the warp index or slice. The Sliced ELLPACK and ELLR-T (many threads per row) ideas are combined in the Sliced

ELLR-T scheme of Dziekonski et al. [37].

A series of papers considered the block compressed sparse row format (BCSR). BCSR stores the non-zero entries as a sequence of fixed-size $r \times c$ dense blocks. BCSR uses one integer index of storage per block instead of one per non-zero as in CSR, reducing the index storage by $1/(rc)$. Moreover, fixed-sized blocks enable unrolling and register-level tiling of each block-multiply. The speedup achieved is unfortunately mitigated by fill-in of explicit zeros resulting from this storage. Vuduc et al. [38] mitigated these issues by splitting the matrix into a small set of matrices, each with its own (r, c) and using a more flexible storage scheme called UBCSR. In [7], the BCSR scheme (called in that paper BCRS — block compressed row storage) is used to implement a sparse general-purpose linear solver (i.e., Jacobi-preconditioned Conjugate Gradient). Monakov [39] explored a variant with a hybrid BCSR/BCOO format. BCSR uses a CSR-like format to store blocks. The column coordinate for each block is stored, and the row coordinate is implicitly encoded by sorting blocks by rows and then storing the index of the first block in each row in a CSR-like format. On the other hand, BCOO uses a blocked coordinate format where both coordinates of a block are stored. This provides additional flexibility in comparison to a pure BCSR format.

We mention a separate line of research by S. Baxter [40]. In the ModernGPU library, Baxter considers a SPMV as a sequence of two-steps. Step 1 is comprised of element-by-element multiplication of matrix entries by vector entries. This step is extremely parallelizable and runs near peak performance on GPUs. Step 2 is comprised of a segmented reduction (i.e., a large sequence of short reductions) to produce the final output vector. The focus is therefore on implementing an efficient segmented reduction. The resulting algorithm has high efficiency but is also very complex to implement and modify. The overall approach is very different compared to other SPMV papers. Nonetheless it has produced some of the fastest SPMV GPU codes currently available.

Our new algorithm follows in the footsteps of the "padded jagged diagonals storage" (pJDS) scheme [41]. This is one of the most efficient storage formats available. It reduces the memory footprint by up to 70%, and achieves 95% to 130% of the ELLPACK-R performance. A main drawback of ELLPACK-R and its sliced version is that when rows are highly inhomogeneous (the number of non-zeros varies significantly), memory and floating-point operations are wasted and this overhead may become notable. In contrast, pJDS sorts all the rows from longest to shortest. The rows are then grouped into slices (as in sliced ELLPACK-R), and then packed with zeros up to the longest row. Because of the sorting, the extra padding is minimal and most floating-point operations become more efficient. Here is the pseudo-code for this algorithm:

Listing 1: Pseudo-code for the pJDS scheme

```
for (i=0; i < N; ++i)
  for (j=0; j < rowmax[i]; ++j){
    col_offset = col_start[j];
    c[i] += val[col_offset + i] * rhs[ col_idx[col_offset + i] ];
  }
}
```

Figure 1: N is the number of rows in the matrix, while the array val holds the elements to the matrix. rhs is the vector to be multiplied in SPMV and col_idx are the column indices pertaining to the elements of val. rowmax designates the length of a particular row. col_offset denotes the start of a particular column of the matrix val.

The permutation on the rows of the matrix means that entries in the vector should be permuted as well. There is a computational cost associated with that. However in the context of iterative methods, this overhead is mitigated. The approach in this paper has a similar drawback.

In developing a GPU-based finite element code for real-time heart simulations, we developed a variant of pJDS. We found that it works very well for finite element unstructured grids as well as a variety of other sparse matrices as found in the benchmark suite of Williams et al. [42, 43]. Our algorithms take a novel approach by systematically addressing issues with current GPU algorithms related to row length irregularity, padding efficiency of matrices, and data localization for GPUs. While other algorithms have similar insights in gaining additional performance for finite element SPMV operations by grouping rows with the same row structure [38, 41], encoding sparse matrices in blocks [7, 36, 37], or utilizing data coalescing repacking algorithms [35, 40], our algorithms were designed to

reduce GPU data locality issues, increase memory throughput, and perform well on relatively unstructured meshes with a substantial degree of irregularity in row lengths distributions.

Compared to the pJDS scheme (row sorting) discussed above, our algorithm differs in a few ways. We use a warp optimized storage. As in pJDS, rows are grouped together into sets (slice) such that a single warp processes a given set. Inside a set, the number of non-zero entries is taken to be constant and is set to be equal to the maximum exact number of non-zeros in the row set. Entries for a given set are stored contiguously in memory such that the memory access by threads is perfectly coalesced. A varying number of threads is assigned to a given row. This allows optimizing situations in which few rows are significantly longer than others. In that case, a large number of threads (or all 32 threads in a warp) can be assigned to that row. As a result, per warp, we basically only need to know 3 integers for memory access: the starting address in the matrix A, the number of rows processed by the warp, and the number of non-zero entries per thread. Because of the permutation on the rows of the matrix (as in pJDS), the sequence of column indices inside a given row is typically "random". We sorted them in order to improve cache performance when accessing the right-hand side vector x. In some sense our algorithm combines insights from ELLPACK-R (varying the number of non-zeros K), sliced format (reduced storage), ELLR-T (many threads per row; in our case the number of threads per row is not even constant and is adjusted for load-balancing), and pJDS (sorting of rows). We called our new scheme ELL-WARP.

During the review process, it came to our attention that an algorithm called SELL-C-σ was proposed in a recent paper submitted for publication and available via ArXiv [44]. Their publication compares Intel, Intel Phi, and Nvidia architectures and they examine different metrics with which to evaluate their SELL-C-σ algorithm. We developed our algorithms separately and concurrently, and our initial algorithm K1 is equivalent to SELL-C-σ. Their paper also highlights difficulties in addressing particular types of matrices which we have addressed fairly well in our K2 algorithm. This paper concentrates more on examining the GPU aspects of implementation, the important details regarding the necessary reordering costs of these classes of algorithms, and also proposes some techniques of ameliorating those issues. Lastly, while this paper confirms some of the observations of the aforementioned paper, we also provide further insight into other important aspects of SPMV computational performance differences by carefully modifying our proposed kernels.

We also note that within the context of the finite element method, our novel SPMV routines can be leveraged both for the global assembly operation and within the solver to gain additional computational performance.

In this paper, we first describe the computational heart simulation problem within the context of a traditional finite element framework. We then highlight computational costs for both the assembler and solver, and highlight differences between CPU and GPU implementations. A survey of different GPU SPMV algorithms then follows. A new set of SPMV algorithms is introduced, and a variety of SPMV algorithms are benchmarked against our novel algorithms. Different aspects of optimizations used in the development of our algorithms are investigated. We discuss the utility of our SPMV algorithms in the contexts of finite elements and general computational settings. Lastly, we end with concluding remarks where we highlight possible future improvements to our algorithms, and also discuss possible improvements in leveraging GPUs for finite element simulations.

2. DESCRIPTION OF THE PHYSICAL PROBLEM

Our original aim was to utilize GPUs for achieving real-time electrophysiological simulations of the heart. The electrical propagation of voltage within the heart is described by the following two-variable phenomenological governing differential equations for nonlinear mono-domain excitable tissues.

$$\dot{\phi} = \operatorname{div} \boldsymbol{q}(\phi) + f^{\phi}(\phi, r) \tag{1}$$

$$\dot{r} = f^{r}(\phi, r) \tag{2}$$

The transmembrane voltage, ϕ, is the voltage difference across the cell membrane, while the phenomenological variable, r, describes the recovery behavior of cardiac tissue. To account for conduction throughout the tissue, a phenomenological potential flux term div $\boldsymbol{q}$ is added.

$$\boldsymbol{q} = \boldsymbol{D} \cdot \nabla \phi \tag{3}$$

A phenomenological diffusion tensor $\boldsymbol{D} = d_{\text{iso}}\boldsymbol{I} + d_{\text{ani}}\boldsymbol{n} \otimes \boldsymbol{n}$ allows for cell-to-cell electrical coupling across cellular gap junctions. d_{iso} and d_{ani} are the respective isotropic and anisotropic conduction terms and $\boldsymbol{n}$ is the preferred direction of anisotropic conduction.

In this paper, we will use the classical Aliev-Panfilov model [45] for convenience to evaluate the effectiveness of our GPU algorithmns. The source terms for the Aliev-Panfilov model are

$$f^{\phi} = c\,\phi\,[\phi - \alpha][1 - \phi] - r\,\phi \tag{4}$$

$$f^{r} = \left[\gamma + \frac{\mu_1\, r}{\mu_2 + \phi}\right]\,[-r - c\,\phi\,[\phi - b - 1]] \tag{5}$$

While the mono-domain equations (1) and (2) can be solved simultaneously, a global-local internal variable splitting approach [46, 47] is taken because it yields a global symmetric tangent matrix for the global degree of freedom, ϕ, such that a fast iterative GPU solver can be effectively employed. In our splitting approach, we take the transmembrane voltage, ϕ, to be our global variable and r to be our local variable. r is local in the sense that the value of r at a given point in time is not directly affect by the current values at the neighboring points. r at a given point is affected by neighboring points through coupling through ϕ.

An implicit semi-discretization approach is taken where backward Euler time integration is used and the mesh is discretized spatially using Lagrangian shape functions. The heart mesh is discretized in space using isoparametric tetrahedral elements in this paper using equations (6) N are tetrahedral shape functions, while $\delta\phi_i$ and ϕ_j are the respective test function local nodal values and local nodal solution values. $\mathcal{B}^e$ is the domain of the element and n_{en} is the number of nodes within a particular element.

$$\delta\phi^h|_{\mathcal{B}^e} = \sum_{i=1}^{n_{\text{en}}} N^i\ \delta\phi_i \qquad \phi^h|_{\mathcal{B}^e} = \sum_{j=1}^{n_{\text{en}}} N^j\ \phi_j \tag{6}$$

After the initial spacial discretization has been performed, initial values are set for the global variables ,ϕ_J, at the finite element nodes, while the internal variable r is initialized at each integration point within each element. As the coupled nonlinear problem is stiff, backward Euler implicit time integration is used where ϕ_n denotes the fully converged solution at the previous timestep.

$$\dot{\phi} \approx \tfrac{1}{\Delta t}[\,\phi - \phi_n\,]$$

For a given timestep, $n+1$, Newton's method is applied over the finite element nodal values, ϕ_J, such that the global residual, R^{ϕ} is fulfilled. In order to ensure that the residual for the internal variable, R^{r} is also fulfilled over the whole domain, we apply Newton's method at every integration point within every element and update r such that R^{r} is fulfilled at every point. Therefore, the nonlinear rate equation (2) is first solved locally using equation 7. The tangent matrix K^{r} and residual R^{r} are constructed at each integration point and used to iteratively determine the converged nonlinear solution for the given timestep , $n+1$.

$$\mathrm{R}^{r} = \dot{r} - f^{r}(\phi, r) \qquad \mathrm{K}^{r} = \mathrm{d}_{r}\mathrm{R}^{r} = \frac{d\mathrm{R}^{r}}{dr} \tag{7}$$

With r updated based on ϕ_J, the current Newton iteration value of the transmembrane potential at the current timestep at each finite element node, the terms for the global residual, R_I^{ϕ} can be calculated using equations 8. The tangent term $\mathrm{d}_{\phi} f^{\phi}(\phi, r)$ necessary for computing $\mathrm{R}_I^{\phi\,\mathrm{e}}$, the global residual contributions from each element, can be computed in a straightforward manner and the tangent term is used to obtain the completed global element tangent matrix $\mathsf{K}_{IJ}^{\phi\,\mathrm{e}}$ for each element. The details for this procedure can be found here [47]. The global element residuals $\mathsf{R}_I^{\phi\,\mathrm{e}}$ are then calculated and together with the element tangent matrices are assembled to form the global tangent matrix K_{IJ}^{ϕ} and global residual vector R_I^{ϕ} , as is done normally in nonlinear finite elements.

$$\mathsf{R}_I^{\phi\,\mathrm{e}} = \dot{\phi}^{\mathrm{e}} - \operatorname{div} \boldsymbol{q}(\phi^{\mathrm{e}}) - f^{\phi}(\phi^{\mathrm{e}}, r^{\mathrm{e}}) \qquad \mathsf{K}_{IJ}^{\phi\,\mathrm{e}} = \mathrm{d}_{\phi_J^{\mathrm{e}}}\mathsf{R}_I^{\phi\,\mathrm{e}} = \frac{d\mathsf{R}_I^{\phi\,\mathrm{e}}}{d\phi_J^{\mathrm{e}}} \tag{8}$$

With the global tangent matrix and residual, the global finite element problem is solved iteratively using Newton's method over the domain. Both the local and global degrees of freedom are updated at each global Newton-Raphson iteration to ensure the consistency of the solution. This process is repeated until convergence is reached at each timestep. The scheme is summarized in Table 1.

Table 1: An algorithmic treatment of the transmembrane voltage in excitable cardiac tissue based on finite element discretization in space and implicit finite difference discretization in time embedded in two nested Newton-Raphson iterations. The electrical unknown, the membrane potential ϕ, is introduced globally on the node point level whereas the phenomenological recovery variable, r, is introduced locally on the integration point level. Element assembly is highlighted in blue, and the solution update is highlighted in red.

initialize nodal degrees of freedom ϕ_J initialize internal variable r
global Newton iteration
loop over all elements
loop over all integration points
local Newton iteration
calculate local recovery variable residual, $\mathsf{R}^r = [r - r^n]/\Delta t - f^r$ and local tangent matrix $[\mathsf{K}^r] = \mathrm{d}_r \mathsf{R}^r$ update the recovery variable $r \leftarrow r - [\mathsf{K}^r]^{-1}\mathsf{R}^r$
calculate source term $f^\phi(\phi, r)$ and its linearization $\mathrm{d}_\phi f^\phi$
calculate element residuals $\mathsf{R}_I^{\phi\,\mathrm{e}}$ and element matrices $\mathsf{K}_{IJ}^{\phi\,\mathrm{e}} = \mathrm{d}_{\phi_J^\mathrm{e}} \mathsf{R}_I^{\phi\,\mathrm{e}}$
calculate global residual R_I^ϕ and global iteration matrix $\mathsf{K}_{IJ}^\phi = \mathrm{d}_{\phi_J} \mathsf{R}_I^\phi$
update membrane potential $\phi_J \leftarrow \phi_J - \mathsf{K}_{IJ}^{\phi\,-1} \mathsf{R}_I^\phi$

Implicit time integration can be further leveraged by using adaptive time-stepping schemes to further reduce the computation time. Lastly reduced-order integration is used, whereby a linear tetrahedral element will only have one integration point. All implementations developed in this study are scientifically accurate and run at double precision.

2.1. Computational Cost of the Finite Element Method

While the biophysical problem may initially seem rather specific, the problem has been structured in the standard non-linear solid mechanics finite element framework. For example, such local-global splitting schemes are commonly used in the field of plasticity [48]. A comparison between the computational cost of our heart simulation for various refinements of the same mesh are shown in Figure 2 for a single CPU finite element method implementation of the electrophysiological problem using a standard scientific code, PETSc [49], and for our custom single GPU implementation.A Jacobi preconditioner was used for both the CPU and GPU solver. Four different mesh refinements were created having roughly 3,000 nodes, 5,000 nodes, 30,000 nodes, and 50,000 nodes. These meshes have 11k, 16k, 123k and 218k elements respectively. Our GPU implementation utilizes sparse matrix vector (SPMV) multiplication operations to assemble the global tangent and residual quantities, and SPMV is also employed in our standard Jacobi preconditioned conjugate gradient solver. The use of SPMV within the context of CG solvers is prevalent [7, 8, 9, 38, 39]; however the benefit in using SPMV for global assembly on GPUs is not particularly obvious and will be explained in more detail in the following sections.

For our GPU heart simulations the portion of time spent on SPMV during assembly is relatively small compared

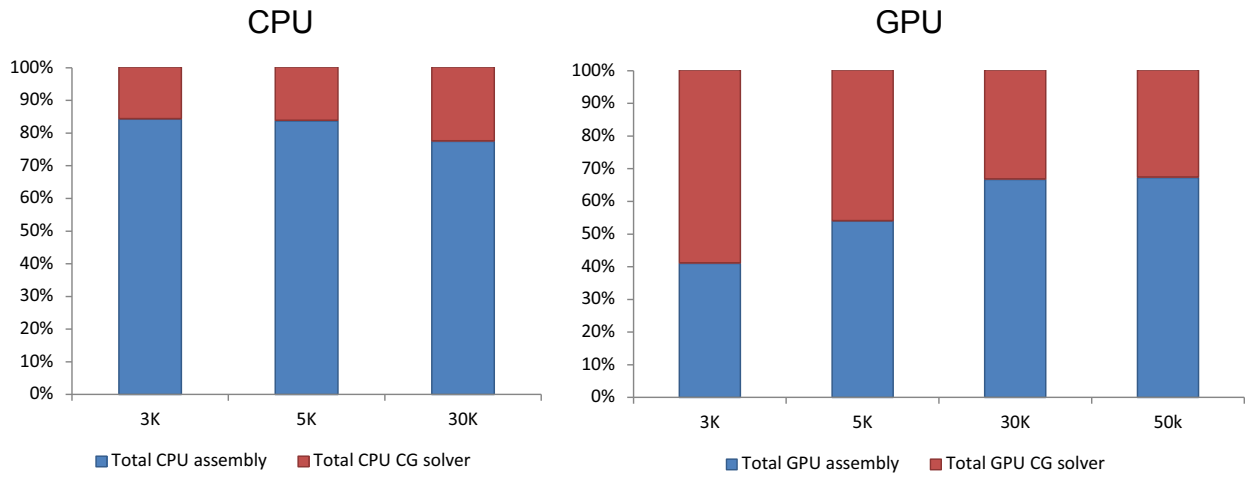

Figure 2: Percentages of computation time of finite element assembly (red) and the CG solver (blue) of the heart problem compared to the total simulation time are shown for both the CPU (left) and GPU (right). Results are shown for the same finite element heart geometry at different refinements: 3K, 5K, 30K, and 50K nodes. A Jacobi preconditioner was used for both the CPU and GPU solver. The PETSc solver using a single core failed to converge for the 50K mesh when the simple Jacobi preconditioner was used.

to the total assembly time and increases towards 33% within the CG solver as the level of refinement increases (Figure 3). Together, the total time spent in our GPU implementation shows that up to 10% of the computation time is spent purely on SPMV calculation. However, if other parts of the code are further optimized and accelerated, it is possible that the SPMV calculation could become a considerable bottleneck. Thus, while general finite element problems may exhibit different solver-to-assembly ratios, improvements to SPMV operations can be beneficial to finite element frameworks. In the following sections, we will investigate various improvements and novel modifications to current SPMV algorithms, and highlight their use within the finite element method and in other general SPMV applications.

3. GPU ARCHITECTURAL ORGANIZATION

As GPUs are architecturally organized differently than traditional CPUs, a brief summary is given before we examine current GPU algorithms. These architectural differences are important in developing a substantially improved SPMV algorithm. Since we are using an Nvidia GPU, we will describe the general organization of an Nvidia GPU. The overall work-execution model on the GPU is composed of individual threads that each execute a specified kernel routine concurrently with respect to other threads (Single-Instruction Multiple-Thread = SIMT). Groups of 32 threads, called warps, are executed synchronously. Furthermore, up to 32 warps can be grouped into a block (CUDA Compute Capability 2.0). Threads from a block have access to shared memory within that block and can also be forced to synchronize with other threads within that block.

On the GPU, there are basically four types of memory available: global, shared, local, and register memory. In general, memory that is available to more threads is slower than memory that is shared within a smaller group of threads. For example, global memory can be accessed by all threads on a GPU, but it is also the slowest form of memory available. Shared memory is shared within a block. It is always faster than global memory. However memory bank conflicts may occur when threads are accessing the same shared memory banks simultaneously. Local memory is global memory reserved for threads where stores are kept in the L1 cache. Generally its use is specified by the compiler and it is used when a thread has fully utilized the maximum amount of memory registers available. Registers are the fastest form of memory, and only 63 registers are available for a given thread for the GPU used in this work. Different memory caches are used to facilitate memory accesses to and from the different

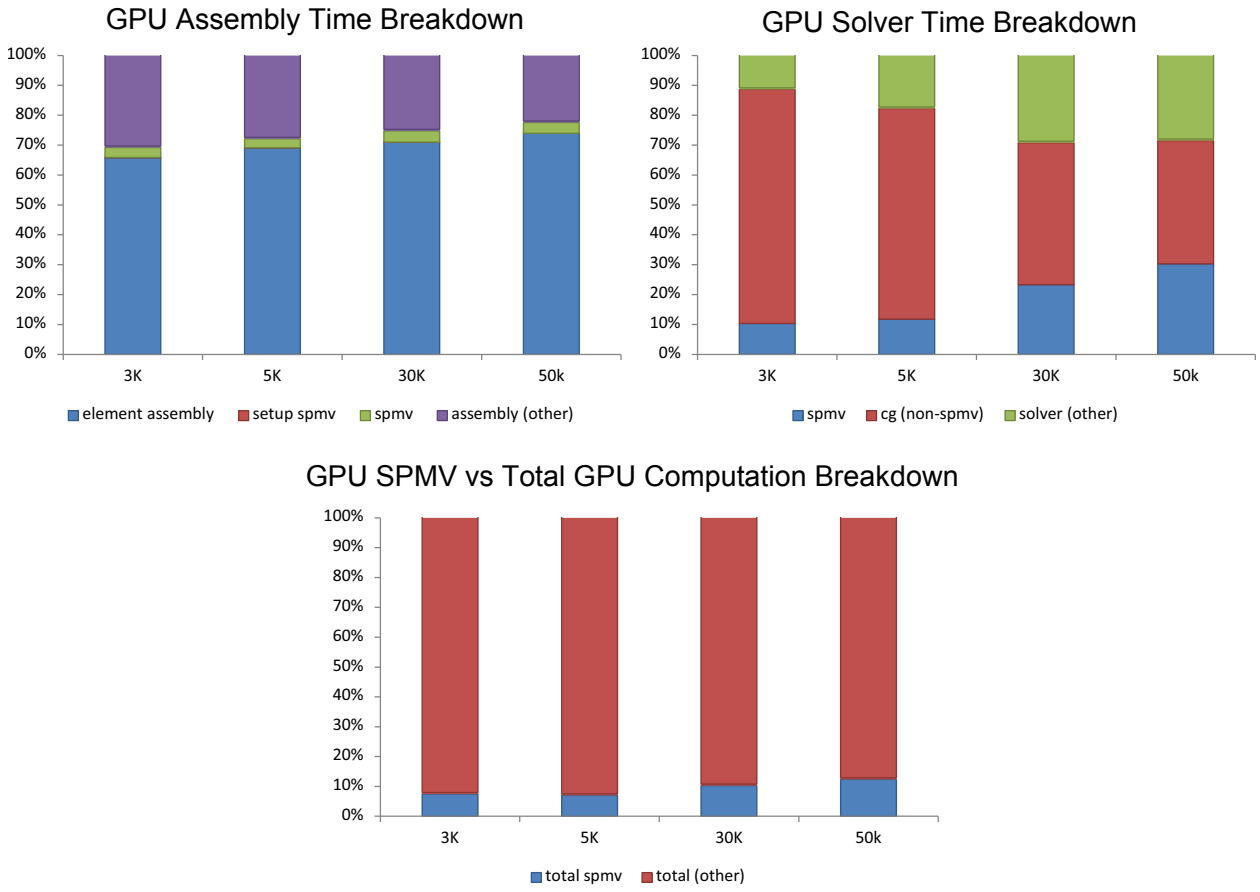

Figure 3: The proportion of time spent on different finite element assembly tasks within our GPU implementation for different refinements is shown (top left). SPMV is used both during the assembly and during the CG solver. Element assembly was performed on a thread level where each thread assembles an element and is shown in blue. The time spent preparing for the SPMV operation is shown in red, while the actual SPMV operation is denoted in green. The time spent initializing, organizing, and copying to and from GPU memory is denoted in purple. Proportions of SPMV with relation to other tasks for the CG solver are shown (top right). The time spent performing the SPMV operation (blue), the time spent performing the CG algorithm excluding time spent performing the SPMV (red), and initialization and GPU memory transfer (green) are indicated for each mesh refinement. The total proportional time spent on SPMV for both assembly and solver portions (blue) in comparison to the rest of the GPU finite element method (red) is shown (bottom). CUSPARSE [33] was used to carry out all SPMV operations in these results.

types of memory to the individual threads. In this paper, the L1 and L2 caches are implicitly used and use of the texture cache, which is a special spatial locality-based global memory cache, is also investigated.

4. DESCRIPTION OF CURRENT GPU METHODS

In order to achieve real-time finite element heart simulations, ideally both the finite element tangent matrix assembly and finite element solver will perform a majority of the computations on the GPU to minimize CPU-to-GPU memory transfer overhead. Current GPU approaches and methods are summarized below.

For the finite element assembly of cardiac simulations, efforts have been made at leveraging GPUs to speed up the evaluation of the local O.D.E. problem (2) [26, 27]. Since we have chosen a simple phenomenological model, these efforts are not directly applicable to this study. However, several different finite element assembly approaches have been recently [20, 21, 50, 51]. In Cecka et al. it is shown that assembly by non-zeros of the finite element tangent matrix is a better approach compared to assembly by elements, which is the traditional serial finite element approach. Furthermore, if we restructure the assembly such that shared memory is used instead of global memory, further speedups are possible. However, this improvement requires a substantial rewrite of the finite element method and is dependent on the particular type of element. The method of coloring, where elements are specified in a way such that threads can safely accumulate entries in the global matrix without potential race conditions, can also be used, but multiple passes are typically required and memory traffic may increase. These consequences may affect the overall performance, however the number of flops is often nearly optimal. Also of note, in Markall et al. 2013, a method for assembly the residual vector is proposed as the best solution for their problem; however they propose a different scheme for the global matrix assembly. Much work has also been done on finite element CG solvers and multi-grid solvers on GPUs [2, 7, 9, 38]. Lastly, both the finite element assembly on the GPU and conjugate gradient method utilize forms of sparse matrix vector product multiplications.

In this paper, we take a hybrid approach where element quantities are first obtained in parallel. However instead of directly assembling the global tangent matrix directly from each finite element's stiffness matrix, these element quantities, $\mathsf{K}^{\phi\,\mathrm{e}}_{IJ}$, are organized into an SPMV problem where each row corresponds to the assembly of a non-zero element in the global matrix, K^{ϕ}_{IJ}. This approach attempts to avoid race conditions and retain efficiency while avoiding large deviations from standard finite element codes. In our approach, the finite element assembly procedure can be reorganized into two separate SPMV operations to build the tangent matrix, K, and residual vector, R, for non-linear iterative finite element simulations. While it is possible to combine the assembly into one operation where the memory associated with K and R is adjacent, for the sake of clarity in this paper we will treat them as separate SPMV operations.

$$\boldsymbol{K}^{\mathrm{elem}} = \begin{bmatrix} \mathsf{K}^{\phi\,1}_{1,1} & \mathsf{K}^{\phi\,2}_{1,1} & \mathsf{K}^{\phi\,3}_{1,1} & \mathsf{K}^{\phi\,4}_{1,1} & \dots \\ \mathsf{K}^{\phi\,1}_{1,2} & \mathsf{K}^{\phi\,4}_{1,2} & 0 & \dots & \\ \vdots & & & & \\ \mathsf{K}^{\phi\,7}_{\mathrm{n_{nodes}},\mathrm{n_{nodes}}} & 0 & \dots & & \end{bmatrix}, \qquad \boldsymbol{K}^{\mathrm{elem}} \cdot \bar{\mathbf{1}} = \mathsf{K}^{\phi}_{\mathrm{flat}}$$

In our problem, $\boldsymbol{K}^{\mathrm{elem}}$ is the temporary sparse tangent matrix used to assemble the global tangent matrix. Each row corresponds to a particular entry, $\mathsf{K}^{\phi}_{\mathrm{IJ}}$. Therefore $\boldsymbol{K}^{\mathrm{elem}}$ has $\mathrm{nnz_K}$ rows, where $\mathrm{nnz_K}$ is the number of nonzero global elements in $\mathbf{K}^{\boldsymbol{\phi}}$. Entries within a row are simply the different contributions from the various element stiffness matrices $\mathsf{K}^{\phi\,\mathrm{e}}_{\mathrm{IJ}}$ associated with the particular tangent matrix location. By summing up each row the global matrix entries, K^{ϕ}_{IJ} can be accumulated without any race conditions in parallel. This, in fact, can be described as a trivial SPMV operation (shown above). $\bar{\mathbf{1}}$ is a vector of 1's of length $\max(n^{\mathrm{i}}_{\mathrm{conn}})$, where $n^{\mathrm{i}}_{\mathrm{conn}}$ is the number of elements that share a node i. The resulting global matrix K^{ϕ} can be flattened into a vector, $\mathsf{K}^{\phi}_{\mathrm{flat}}$, that has $\mathrm{nnz_K}$ entries. In the same way, the residual vector entries, $\boldsymbol{R}^{\mathrm{elem}}$, can be assembled using the aforementioned SPMV organization used for assembling R^{ϕ}. $\boldsymbol{R}^{\mathrm{elem}}$ is a matrix of size $n_{\mathrm{nodes}} \times \max(n^{\mathrm{i}}_{\mathrm{conn}})$. We note that it is trivial to modify a SPMV algorithm such that only summations are performed without needing to multiply each elemental quantity by 1.

The CG solver is then used to solve the canonical problem $\mathbf{K}^{\boldsymbol{\phi}}\,\Delta\boldsymbol{x} = \mathbf{R}^{\boldsymbol{\phi}}$, where $\Delta\boldsymbol{x}$ is the solution update for the Newton-Raphson method. At least one SPMV is performed during each CG iteration step. Therefore, by

improving SPMV algorithms on GPUs, we can improve different parts of the finite element method and also apply these algorithms to general SPMV linear algebra problems.

While there are many variations of sparse matrix vector product GPU algorithms [39, 7, 38], the majority can be summarized by the available implementations from several commonly available sparse matrix vector libraries: CUSPARSE [33], CUSP-library [32], and ModernGPU [40]. In [29], COO, CSR, ELL, and HYB formats are examined and analyzed. A review of the different formats and algorithms is highlighted below. In this section, we will look at the classical SPMV problem $\boldsymbol{A}\,\boldsymbol{x} = \boldsymbol{y}$, where $\boldsymbol{x}$ is given and the objective is to calculate the vector $\boldsymbol{y}$.

COO

The coordinate list (COO) sparse matrix format is composed of three array lists of row indices, column indices, and the corresponding list of matrix values. In the CUSP COO implementation, row indices are sorted in order. Each warp processes a section of the matrix and works over 32 non-zero values at a time. Each thread within the warp performs a multiplication between the thread's value and the corresponding value from the vector. The different row segments are summed using segmented reduction, which is a method of efficiently distributing the reduction over a warp even if the warp is computing the sum for different rows. During the initial stage, the product is stored in shared memory and intra-warp segmented reduction is performed on the element products. The first thread within each warp determines whether to include results from a previous warp iteration in its row sum, and if not, updates the solution vector. Intra-block segmented reduction is then used to properly accumulate rows that span multiple warps.

COO SPMV algorithms generally suffer from poor memory to computation ratios as row and column indices must be retrieved for each computation. Row indices are also required for row sum reduction, and additional explicit intra-block thread synchronization may be required depending if rows span multiple warps.

CSR

The compressed sparse row (CSR) matrix storage format is composed of three array lists of values, corresponding column indices, and corresponding row offsets that index the beginning of each row in the arrays of values. In the CUSP CSR vector implementation, a warp is assigned to each row in the matrix. A warp processes a continuous section of values and corresponding column indices in a coalesced manner. Intra-warp reduction is performed and sums are accumulated in the solution vector.

The main issue with CSR algorithms is that while the storage is space efficient and the memory access pattern is contiguous, memory accesses are not aligned. The CUSP implementation will also suffer when the size of a row is less than the warp size. However, unlike an alternative CSR algorithm where each thread is assigned to a row, the access pattern in CUSP is coalesced and contiguous.

ELL and ELL variants

In the ELLPACK (ELL) format, the matrix is again described by two array lists of values and column indices. The maximum size of the longest row in the matrix is allocated for each row in the ELL format. Non-zero values are ordered contiguously, and the remaining values are typically zero-padded. Column indices are arranged and padded in a similar manner. Both values and columns are stored in column-major order. The GPU kernel is fairly simple. Each thread is responsible for a row, and matrix vector products are summed in a coalesced manner.

The ELL format performs poorly when the row size varies from row to row resulting in poor work distribution within each warp. The format also suffers from issues with excess padding when the longest and shortest row differ greatly in length, and when the average row length is significantly shorter than the longest row.

There are several variations of the ELL format which attempt to fix these issues. One is the ELL-R format [34] where an extra vector is provided that designates the number of non-zeros each row is responsible for. Another is Sliced ELLPACK [36] where the matrix is sliced into groups of consecutive rows and then stored in the ELLPACK format. The ELLR-T [35] and sliced ELLR-T formats [37] use multiple threads per row. Lastly, another variant is the padded jagged diagonals storage (pJDS) format [41] which is similar to ELL-R. In this format, the rows are first sorted from the longest to shortest row, and then sliced and padded in a manner similar to the ELL-R format.

HYB

The Hybrid format is similar to the ELL format. To reduce the amount of zero-padding due to row-to-row size differences, a certain number of values, determined empirically, is stored in ELL, and the remaining entries are stored in COO format. The CUSP implementation is simply a combination of the ELL kernel and the COO kernel.

Due to the more efficient storage scheme, the hybrid format requires launching of two kernels in order to perform the sum. While the values of the matrix are contiguous in each kernel, the access pattern for the vector will generally not be so.

ModernGPU

ModernGPU takes a different approach with its sparse matrix vector implementation called mgpusparse (MGPU). The algorithm is somewhat complicated but is essentially constructed to address several issues with the algorithms above. It is a combination of a series of parallel and segmented scans. Data is partitioned into the list of values, list of corresponding column indices, and several lists are used to keep track of the ends of rows and row segments. Data is also arranged in a coalesced way similar to column-major storage, but for fixed-sized chunks of matrix values. Two kernels are utilized. In the first kernel, each thread within a warp is responsible for calculating `valuesPerThread` partial row sum entries. A serial scan is performed by each thread and the sums of different row segments are stored in shared memory. A number of flags are associated with each thread, such that each thread can determine where to store the partial sum for a given row segment in shared memory. An intra-warp parallel scan is then performed to reduce rows within each chunk quickly. Lastly an intra-block segmented reduction is performed to reduce rows that span multiple blocks.

MGPU also takes advantage of index representation compression and performs other compression tricks to increase efficiency. The algorithm generally performs well in comparison to the algorithms above. The main drawback, however, is the complexity of the algorithm and the amount of book-keeping that is necessary. The main benefits are that the workload for threads within a block are very well distributed and very few threads are idle compared to the rest of the threads within each block.

5. DEVELOPMENT OF NOVEL GPU ALGORITHMS

There are several other variations of the previously mentioned algorithms. Most variations concentrate on increasing padding efficiency, attaining less warp-divergence, and increasing the workload for a given block of threads [7, 38]. However, there are several key insights that have been made in ModernGPU, ELLPACK, and CSR vector implementations that may yield a better SPMV algorithm for finite elements. In ELL and MGPU algorithms, "column-major" coalesced data interleaving is leveraged to allow threads within a warp to work on different rows and row-segments. The interleaving results in an efficient coalesced memory access pattern. In MGPU and CSR vector, intra-warp reductions are utilized to avoid unnecessary thread synchronization, which allow warps and blocks to execute without delays due to synchronization.

The variations in row length in ELL are acquiesced in the HYB matrix implementation, and are accounted for by segmented scan flags in MGPU. Both algorithms try to address the padding inefficiency in the simple ELL algorithm. In [29], it is noted that the impact of padding on well-behaved structured meshes should not be substantial as the row size should not vary excessively. In unstructured finite element meshes of solids, this appears to be partially true. Since the node connectivity of the mesh directly translates into the finite element tangent matrix and element quality is usually controlled by meshing algorithms, the number of non-zeros per row is indirectly controlled by mesh quality algorithms. Therefore, the row lengths in a given unstructured mesh are usually constrained; however, there is no guarantee that there is row-to-row uniformity within the matrix for a given unstructured finite element mesh. In the case of assembly, the $\boldsymbol{x}$ vector is simply a vector of ones, and one can simply avoiding the multiplication of unity to the corresponding matrix value altogether. From Figure 4, it is clear that while row sizes only vary from 5 to 21 non-zeros per row, the actual distribution of non-zeros per row varies quite a bit from the median and mean row lengths. This distribution poses problems for ELL because the row size varies from row-to-row and therefore the throughput may not be optimal. The distribution also does not yield optimal throughput for CUSP CSR since row sizes are smaller than the size of a warp.

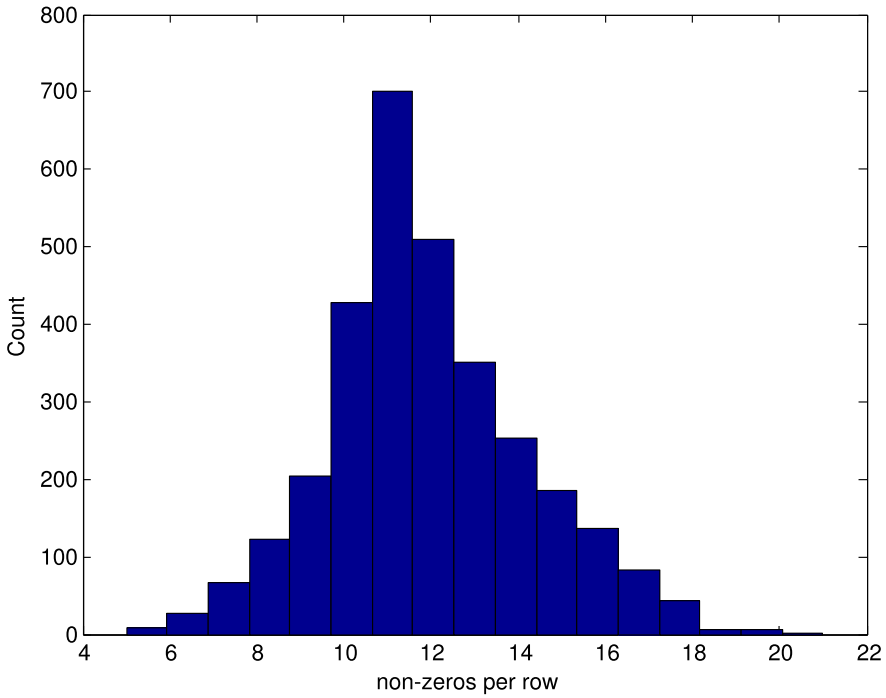

Figure 4: A histogram of the non-zeros per row in a small 3129 row stiffness matrix for a patient-specific simulation of the heart (Heart 3K).

Since MGPU outperforms the simple kernel implementations of ELL/HYB, this seems to indicate that there may be inefficiencies related to padding which are not an issue for the more complicated equal work distributed MGPU implementation. Therefore, we propose the following set of simple kernel algorithms that are directly motivated from the insights listed above.During the preparation of this manuscript, it came to our attention that other ELL variations — sliced ELLR-T and pJDS — share many features with our proposed kernels. While we came up with our algorithms independently, we have made similar key observations. We aim to show how our algorithms are a generalized superset of the ELL variations.

5.1. ELL-WARP

The proposed kernel algorithm is based on ELL. To mitigate the observation of row length distribution irregularity, we first sort the rows by length from longest to shortest (Figure 5). A thread per row execution scheme is used,

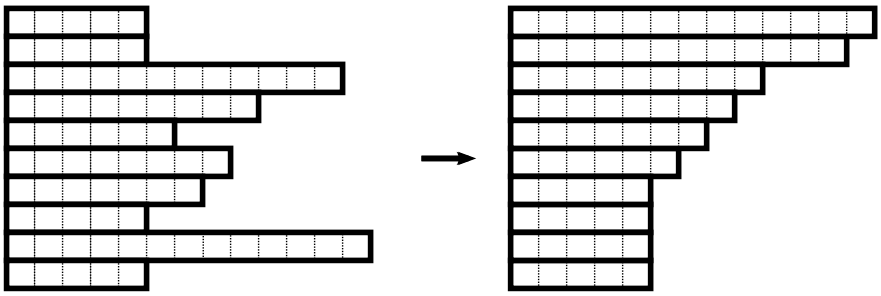

Figure 5: Sort rows of the matrix from longest to shortest.

however, instead of using the representation used in the ELL format where padding continues until the maximum matrix row length is reached, each warp is now padded up to the maximum row length of given warp (Figure 6). To ensure coalesced memory access, we permute the data within each warp in a "column-major" ordering similar

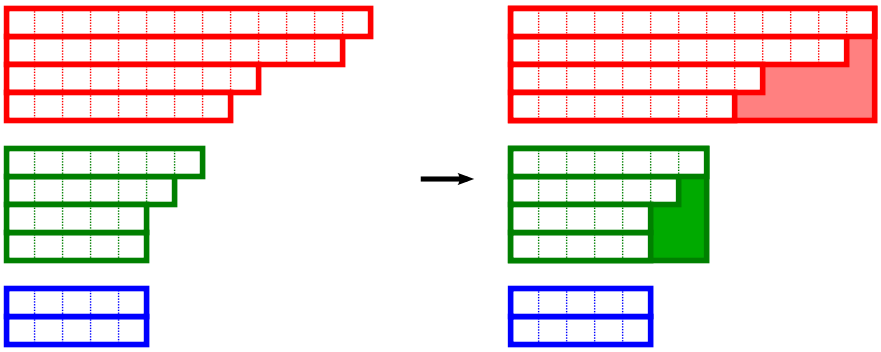

Figure 6: Arrange rows into groups of warp size and then pad accordingly. In this figure, we use a warp size = 4 only for purposes of illustration.

to MGPU and ELL/HYB (Figure 7). Lastly, the results are mapped back to the original row-ordering when writing

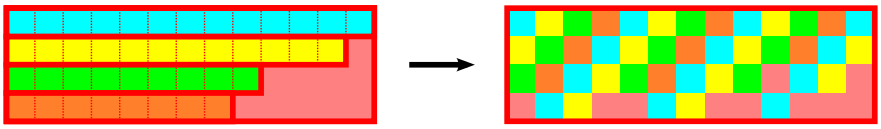

Figure 7: Reorder the first warp in Figure 6 in a coalesced column-major ordering.

results to the solution vector $\boldsymbol{y}$. This algorithm will be referred to as ELL-WARP (K1). The code is shown in Listing 2. K1 requires that a list of values and their column indices be padded per warp, which are then sorted by row length and arranged in "column-major" order. A list of the initial offsets for each warp is given as well as a row mapping from the sorted row index to the original row index.

The purpose of this arrangement is to exploit the relatively small variances in row lengths per warp to reduce unnecessary padding in ELL and also to increase the amount of threads performing meaningful operations. As values and column indices are zero-padded, those values and column indices are cached and should not reduce the effectiveness of this algorithm if the variances within each warp are in fact small. By having a fixed row length for each warp, one can reduce the amount of data needed to compute row sums in comparison to CSR and COO while reducing the number of trivial instructions executed for padded entries in comparison to ELL. Thus, this kernel aims to achieve an equally-distributed workload similar to MGPU in cases where there is some degree of row length regularity within each warp. Unfortunately, the final non-coalesced memory write to the solution vector is a potential inefficiency for this kernel which is otherwise coalesced in terms of memory access to values and column indices.

5.2. ELL-WARP with row reordering (K1r)

While reordering the rows of the matrix helps us achieve row length regularity within individual warps, a straightforward SPMV operation with the row permuted matrix will yield a permuted solution.

$$A_{i,j}x_j = y_i \qquad A_{p(i),j}x_j = y_{p(i)} \tag{9}$$

Thus the final memory store operation undoes this permutation, $p(i)$, such that the original ordering of the solution is preserved; however this operation results in a non-coalesced memory storage pattern. To mitigate the necessity of performing a non-coalesced write within our kernel, the column indices of the matrix and the vector $\boldsymbol{x}$ can be

Listing 2: Code for K1.

```
template <uint WARP_SIZE,bool usecache>
__global__ void K1( *A, *colinds, *Pinv, maxrows, *x, *warp_offset, *y, nrows) {

  const uint tid = threadIdx.x;
  const uint id = tid  + blockIdx.x * blockDim.x;
  const uint wid = tid & (WARP_SIZE-1);
  const uint warpid = id / WARP_SIZE;

  if (id < nrows) {

    IndexType toffset = warp_offset[warpid] + wid;
    uint maxnz = maxrows[warpid] * WARP_SIZE + toffset;

    // Perform sequential sum
    ValueType sum = A[toffset] * cache<usecache> (colinds[toffset],x);
    for(toffset += WARP_SIZE; toffset < maxnz; toffset += WARP_SIZE) {
      sum += A[toffset] * cache<usecache> (colinds[toffset],x);
    }

    // Store remapped result
    y[Pinv[id]] = sum;
  }
}
```

Figure 8: A is an array that contains the entries to the matrix. colinds are the index of the columns. Pinv is the inverse permutation array. x is the x vector. maxrows is an array which contains the maximum row length of a particular warp. warp_offset determines the index at which the first thread of a warp should load from A and colinds. y is the solution vector. nrows is the total number of rows in the matrix. tid is the thread index. id is the global row id corresponding to a particular thread. wid is the index of a thread within a warp. warpid is the global index of a particular warp. toffset is the offset with respect to warp_offset that each thread should load values from A and colinds. maxnz is the used to properly bound the workset of a particular thread. WARP_SIZE and usecache are template variables. WARP_SIZE sets the size of a warp for the given GPU architecture, while usecache determines whether the texture cache is used or not in loading from the x vector.

reordered such that the numbering of the corresponding entries in $\boldsymbol{x}$ and the solution vector are consistent. This is commonly performed in finite element computations by renumbering the unknowns during the setup phase of most codes. Reordering the columns of the matrix and entries of $\boldsymbol{x}$ does not change the result of the solution vector.

$$A_{i,j}x_j = y_i \qquad A_{i,f(j)}x_{f(j)} = y_i \tag{10}$$

By performing this reordering, we can essentially ignore the original row mapping necessary in K1 as the sorted solution vector is now numbered consistently with the newly arranged $\boldsymbol{x}'$ vector. Unfortunately, the reordering of the column indices now produces a non-ordered access pattern when reading values from the $\boldsymbol{x}'$ vector. While the effects may be mitigated by using a texture cache, this is somewhat undesirable.

To further illustrate the reordering scheme, we considered a CSR ordered matrix for convenience. If x^{num} is the original ordering for $\boldsymbol{x}$, then $\boldsymbol{P}$ is the new ordering such that the rows will be sorted accordingly from longest to shortest.

x^{num}	$\begin{bmatrix}1 & 2 & 3 & 4 & 5 & 6 & 7\end{bmatrix}$
$\boldsymbol{P}$	$\begin{bmatrix}2 & 5 & 7 & 3 & 1 & 4 & 6\end{bmatrix}$

The reordered vector, $\boldsymbol{x}'$, and the reordered column indices, $\boldsymbol{c}'$, are defined below, where $\boldsymbol{P}$ is a permutation operator:

$$\boldsymbol{x}'(\cdot) = \boldsymbol{x}(\boldsymbol{P}(\cdot)) \tag{11}$$
$$\boldsymbol{c}'(\cdot) = \boldsymbol{c}(\boldsymbol{P}^{-1}(\cdot)) \tag{12}$$
$$(\cdot) = \boldsymbol{P}(\boldsymbol{P}^{-1}(\cdot)) \tag{13}$$

Given the following single-row CSR matrix $\boldsymbol{A}$, a dense vector of column indices $\boldsymbol{c}$, and a dense vector $\boldsymbol{x}$

$$\boldsymbol{A} = \begin{bmatrix}7 & 8 & 9 & 10 & 2\end{bmatrix}$$
$$\boldsymbol{c} = \begin{bmatrix}1 & 2 & 4 & 5 & 6\end{bmatrix}$$
$$\boldsymbol{x} = \begin{bmatrix}1 & 2 & 3 & 4 & 5 & 6 & 7\end{bmatrix}$$

the K1r scheme described above will re-arrange the data such that the following results.

$$\boldsymbol{c}' = \begin{bmatrix}5 & 1 & 6 & 2 & 7\end{bmatrix}$$
$$\boldsymbol{x}' = \begin{bmatrix}2 & 5 & 7 & 3 & 1 & 4 & 6\end{bmatrix}$$

This reordering scheme does not require that $\boldsymbol{A}$ be ordered, but results in non-ordered column indices, $\boldsymbol{c}'$, which is possibly detrimental to maintaining ordered access from $\boldsymbol{x}'$.

5.3. K1r with sorted column index numbering (K1rs)

The following is a variation on K1r that attempts to fix the issue with non-ordered access to the $\boldsymbol{x}'$ vector. The solution is to renumber the column indices of the matrix A. This requires a reordering of the values of the matrix in each row such that the reordering of the column indices is sorted in order after the reordering of the vector $\boldsymbol{x}$.

Again using the previous single-row CSR matrix as an example, we sort $\boldsymbol{c}'$ such that it is ordered from the smallest to largest indices within each row.

$$\boldsymbol{c}'' = \begin{bmatrix}1 & 2 & 5 & 6 & 7\end{bmatrix}$$

$$\boldsymbol{A}'' = \begin{bmatrix}8 & 10 & 7 & 9 & 2\end{bmatrix}$$

The vector $\boldsymbol{x}'$ remains the same; however, now it is necessary to sort the values of $\boldsymbol{A}$ within each row to properly account for the changes in $\boldsymbol{c}'$. The result is a intra-row reordered CSR matrix, $\boldsymbol{A}''$, and an ordered set of corresponding column-indices, $\boldsymbol{c}''$.

5.4. ELL-WARP v2 (K2)

To accommodate the possibility of larger intra-warp row length variances, a multiple threads per row variation of K1 is proposed. For example, if there exists a small number of abnormally long rows, it may be advantageous to process those rows with more threads while allowing the rest of the rows in the matrix to be processed by K1.

A threshold value for the maximum number of entries a thread should be assigned is prescribed. If a row does not meet this criterion, the row is subdivided in half recursively until it meets this requirement. However, to avoid explicit thread synchronization within a block, if a row is longer than 32 $\times$ threshold, one warp will at the least process the entire row to maintain warp-to-warp row independence. The subdivision and repacking of the matrix is shown graphically in Figure 9.

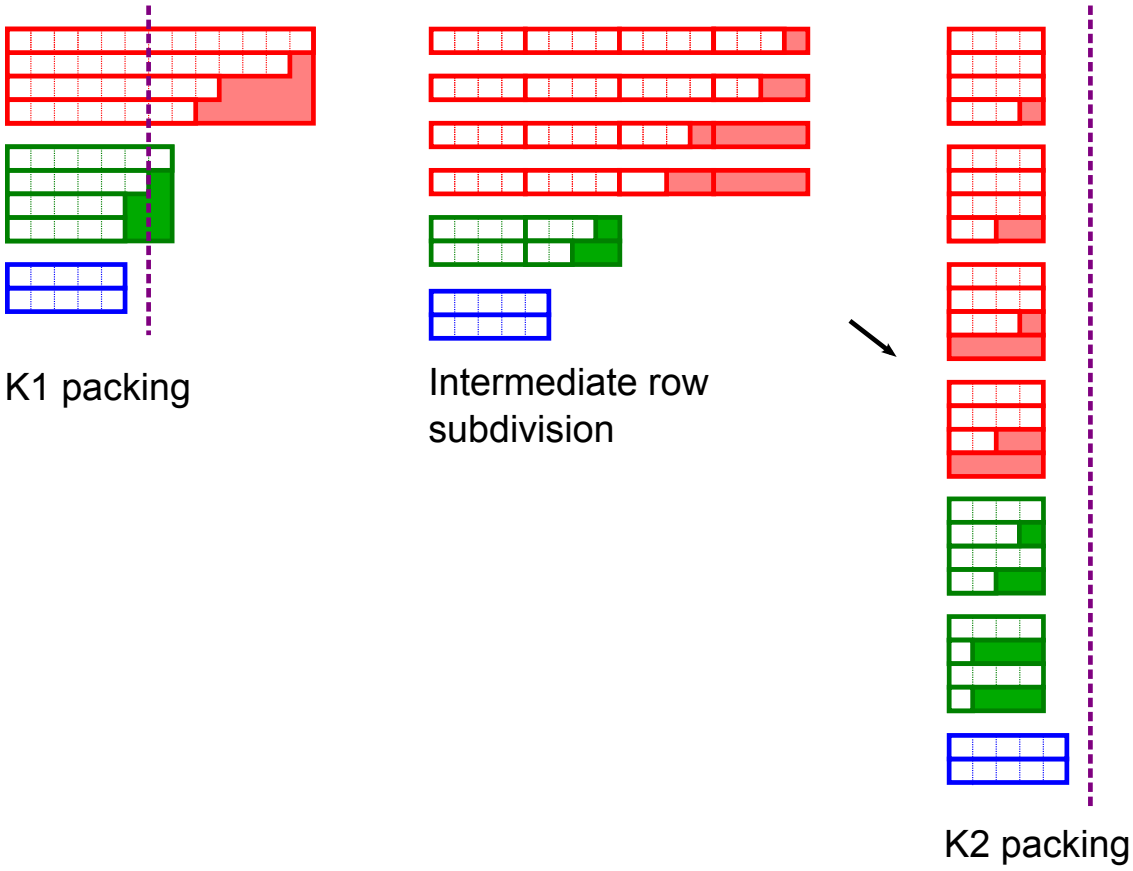

Figure 9: A threshold value (purple dashed line) is given and the original padding in K1 is reorganized into smaller more evenly distributed warps.

The variation inevitably increases the amount of padding in comparison to K1 based on the threshold value chosen. However, it allows unusually long rows to be processed more efficiently, while smaller rows retain the same padding as K1. The threshold parameter forces almost all warps to process the same number of non-zero entries. The kernel now includes a simple parallel reduction loop, which reduces rows in a given warp efficiently using shared memory. This is seen in Figure 10 for the first row of the representative warp in our illustrative example.

As rows are sorted initially, it is relatively easy to reprocess the original warp-sized row groupings from K1 and subdivide the numbers of rows per warp as needed to meet the threshold criterion. Extra information per warp is necessary to dictate the row number per warp and also the number of reductions needed per warp. However, this extra indexing information is minimal and results in two extra integers retrieved from global memory per warp.

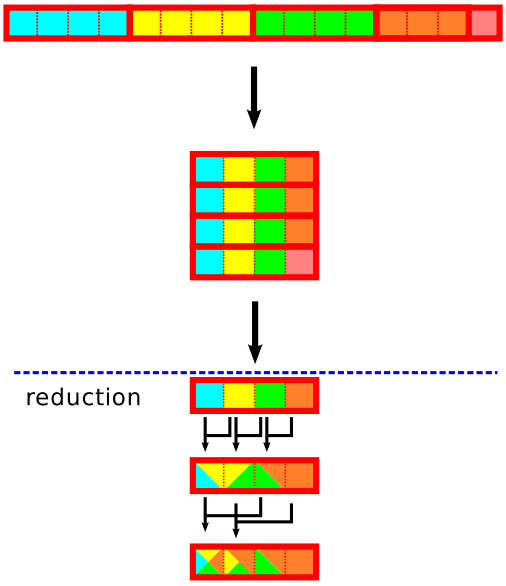

Figure 10: Coloring represents the data for a particular thread. In this instance, 4 threads are used to process the representative row. Data is arranged in column-major order, and then a parallel reduction step occurs resulting with the final sum in the first entry of the output row.

The code for ELL-WARP v2 (K2) is shown in Listing 3.

5.5. K2 variations (K2r) and (K2rs)

Lastly, to avoid the final row mapping from sorted to non-sorted rows, the same variations proposed for K1 can be simply applied directly to K2. The fundamental difference between K1 and K2 is in how the matrix is repackaged into blocks of data. In K2, the packaging is more regular within blocks, whereas in K1, the packaging is more regular only within a given warp.

6. BENCHMARKS OF SPARSE-MATRIX VECTOR MULTIPLICATION METHODS

Our novel algorithms above are compared against several available standard SPMV algorithms using a set of sparse matrices commonly used in benchmarking SPMV methods [42, 43, 29]. Since the K1 and K2 kernels were developed for finite element simulations of the heart, we have also included 3 different refinements of a patient-specific heart mesh as part of the set of benchmark matrices. General information about each of the benchmark matrices is included in Table 2. Matrix structure properties vary substantially from matrix to matrix. It is apparent that certain matrices, such as Circuit and Webbase, have a very skewed non-zero row length distribution. On the other hand, some matrices such as Epidemiology and QCD have very regular row length distributions.

The following kernels were then chosen for comparison: CUSP-CSR and CUSP-HYB [32], CUSPARSE [33], MGPU [40], K1/r/rs, and K2/r/rs. The published MGPU results report the kernel time by performing a large number of products $\boldsymbol{A}\,\boldsymbol{x}$ within a loop, and reporting the average runtime per product. It was observed that the performance of the MGPU kernel varies depending on the number of iterations, with better performance at a higher number of iterations. However, in practice, a single product $\boldsymbol{A}\,\boldsymbol{x}$ is computed followed for some vector operations, after which further products $\boldsymbol{A}\,\boldsymbol{x}$ are needed. Therefore it is unclear to what extent the performance of the MGPU kernel reported for a large number of products $\boldsymbol{A}\,\boldsymbol{x}$, computed immediately after one another, is relevant. Therefore, we will show results for one run of the kernels here, while the results for 1200 iterations can be found in the appendix (Section 11).

Listing 3: Code for K2

```
template <uint WARP_SIZE,bool usecache>
__global__
void K2(ValueType* A, IndexType *colinds,
                    IndexType *rowmap, uint* maxrows,
                    ValueType* x, IndexType *warp_offset,
                    ValueType* y, IndexType nrows,
                    uint* reduction, uint* rows_offset_warp,
                    int nwarps) {

  const uint tid = threadIdx.x;
  const uint id = tid  + blockIdx.x * blockDim.x;
  const uint wid = tid & (WARP_SIZE-1);
  const uint warpid = id / WARP_SIZE;

  extern volatile __shared__ ValueType sumvalues[];

  if (warpid >= nwarps) return;
  const uint offsets = reduction[warpid];
  const uint row_start = rows_offset_warp[warpid];
  const uint rowid = row_start + wid/offsets;

  if (rowid < nrows) {

    IndexType toffset = warp_offset[warpid] + wid;
    const uint maxnz = maxrows[warpid] * WARP_SIZE + toffset;
    ValueType sum = A[toffset] * cache<usecache> (colinds[toffset],x);

    for(toffset += WARP_SIZE; toffset<maxnz; toffset += WARP_SIZE) {
      sum += A[toffset] * cache<usecache> (colinds[toffset],x);
    }
    sumvalues[tid] = sum;

    for (int i = 1; i< offsets; i <<= 1) {
      if (offsets > i ) {
        sum += sumvalues[tid+i];
        sumvalues[tid] = sum;
      }
    }

    if ((wid & (offsets-1)) == 0) {
      y[rowmap[rowid]] = sum;
    }
  }
}
```

Figure 11: Most of the pseudo-code terms are explained in Listing 2. reduction determines how many stages of parallel reduction should be used for a given warp. An offset of 1 means there are 0 stages of reduction, while an offset of 32 means that there are 5 stages of reduction and all threads in a warp are working on one row. rows_offset_warp determines the starting global row index of a warp. row_id is the global row index a particular thread will contribute results to. nwarps is the number of warps needed to compute the SPMV.

Table 2: Matrix benchmark information. The number of nonzeros (nz), the number of rows (nrows), bytes, minimum row length (minrow), maximum row length (maxrow), and average number of non-zero values per row (nz/nrows) is reported for each benchmark matrix. The matrices highlighted (**bold**) are finite element meshes of different refinements of a patient-specific heart mesh.

MatrixName	nz	nrows	bytes	minrow	maxrow	nz/nrows
Circuit	958,936	170,998	19,178,720	1	353	6
Economics	1,273,389	206,500	25,467,780	1	44	7
Epidemiology	2,100,225	525,825	42,004,500	2	4	4
FEMAccelerator	2,624,331	121,192	52,486,620	8	81	22
FEMCantilever	4,007,383	62,451	80,147,660	1	78	65
FEMHarbor	2,374,001	46,835	47,480,020	4	145	51
FEMShip	7,813,404	140,874	156,268,080	24	102	56
FEMSpheres	6,010,480	83,334	120,209,600	1	81	73
Heart 3K	**37,035**	**3,129**	**740,700**	**5**	**21**	**12**
Heart 5K	**52,715**	**4,563**	**1,054,300**	**6**	**22**	**12**
Heart 30K	**367,443**	**28,639**	**7,348,860**	**6**	**24**	**13**
Protein	4,344,765	36,417	86,895,300	18	204	120
QCD	1,916,928	49,152	38,338,560	39	39	39
Webbase	3,105,536	1,000,005	62,110,720	1	4700	4
WindTunnel	11,634,424	217,918	232,688,480	2	180	54

For each matrix used in the benchmark, kernel times are averaged over a specified number of iteration runs. For the K1/r/rs kernels, different block sizes are varied from 32 to 256 threads by increments of 32. Likewise the same is done for the K2/r/rs kernels except, in addition, different maximum non-zero thresholds are varied from the minimum row length of the matrix to the maximum row length of the matrix. Lastly for MGPU, the following number of values per thread are used: 4, 6, 8, 10, 12, and 16. After the parameters for K1, K2, and MGPU are acquired, the fastest averaged times are found. The effective bandwith is reported in Figure 12 and is defined as the rate of bytes processed from the matrix data by a kernel over time.

All benchmarks were run on a single Asus ENGTX480 graphics card (CUDA Compute compatibility 2.0) and on a PC with an Intel I7 950 CPU and 12 GB of memory. Kernels are compiled with optimizations enabled (–O3) and kernels use the texture cache when possible. The results for a single iteration run is shown in Figure 12. The parameters for MGPU, K1 and K2 are shown for each of the best performing kernels for each matrix in Table 3.

The results for a single iteration indicate that our K1 and K2 algorithms perform well overall. For our patient-specific heart meshes, K1 and K2 algorithms perform better than the other tested algorithms. MGPU performs best for 2 matrices: Economics, and FEMCantilever. Meanwhile, K1 and K2 algorithms perform well even for Economics and FEMCantilever, and in general provide the best performance over the benchmarked matrices. Lastly, the effective bandwidth results vary slightly with the number of iterations for MGPU and CUSPARSE, where an increase in the number of iterations improves performance. These results can be found in the Appendix 11. Despite these variances, the K1 and K2 kernels still outperform or are at least comparable to MGPU.

In many ways it is quite surprising that a relatively simple kernel algorithm (22 lines) has comparable performance to a sophisticated segmented scan algorithm in the case of MGPU (145 lines). Unlike other GPU SPMV algorithms, K1 and K2 are monolithic kernels, and extra kernel invocations are unnecessary. However, the thread block size for K1 and K2 are generally different, and thus their parameters must be found individually.

The effect of warp-level organization vs. global data restructuring in HYB can be inferred from Figure 12 since K1 and K2 are partially warp-based variations of the ELL SPMV implementation and thus related to the HYB format. The results from the two variations of ELL show drastically different performance results for the majority of the benchmarked sparse matrices. K1 and K2 reduce the amount of padding with respect to ELL and simultaneously reduce the amount of memory transactions needed for computing the SPMV operation in

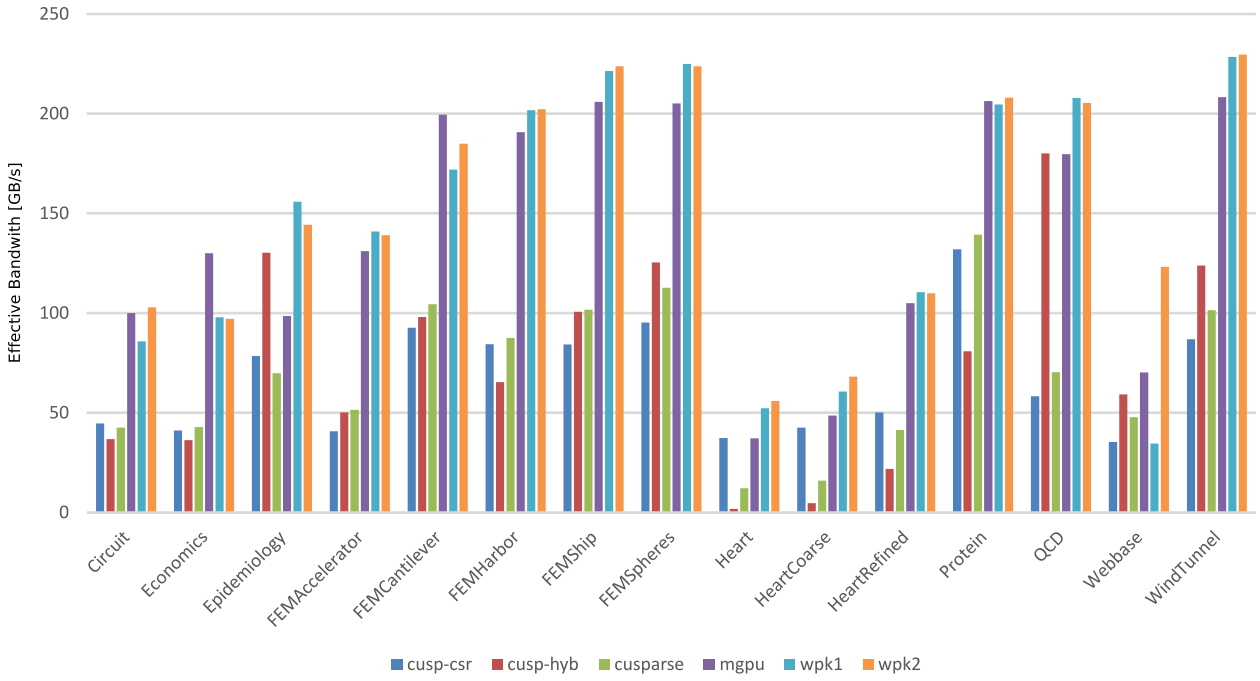

Figure 12: Effective bandwidth benchmark results over matrices for 1 run are shown for CUSP-CSR (blue), CUSP-HYB (red), CUSPARSE (green), MGPU (purple), K1 (cyan), and K2 (orange).

Table 3: Benchmark best parameters for MGPU, K1 and K2.

MatrixName	MGPU	K1	K2	
	valuesPerThread	blocksize	threshold	blocksize
Circuit	6	64	80	160
Economics	6	64	7	96
Epidemiology	6	128	4	128
FEMAccelerator	6	64	58	64
FEMCantilever	10	64	19	224
FEMHarbor	10	96	128	96
FEMShip	10	96	28	128
FEMSpheres	10	96	47	128
Heart 3K	12	32	7	256
Heart 5K	10	64	10	192
Heart 30K	10	64	21	64
Protein	10	64	172	96
QCD	10	224	39	96
Webbase	6	32	658	128
WindTunnel	10	96	38	96

comparison to HYB, thereby providing a dramatic increase in performance. K2 is used when there is a large difference in row length regularity and effectively handles outlier rows in a hybrid ELL-CSR like manner. Together, K1 and K2 are substantially better than ELL and HYB for sparse matrices.

Cost of Reordering Matrix Values

To determine whether it is beneficial overall to reorder the data in a "column-major" coalesced pattern in SPMV applications for finite elements, we consider the following. If K1 and K2 are used to calculate sparse matrix vector multiplications in the conjugate gradient method, only one initial transpose of values is necessary at the beginning of each Newton-Raphson iteration with the assumption that the assembler passes a CSR formatted matrix to the solver. Column indices do not need to be reordered, as we assume that the connectivity of the Lagrangian mesh does not change during the simulation; therefore, only the values of the tangent matrix change while the matrix structure remains constant. We can then determine the number of CG iterations necessary, such that K1 and K2 will outperform CUSPARSE by the following

$$t_{\text{reorder}} + \alpha t_{\text{wpk}} \leq \alpha t_{\text{cusparse}} \tag{14}$$

where α is the number of iterations necessary such that a reordering of data from CSR to "column-major" provides a benefit for the CG solver. t_{reorder} is the time to reorder the CSR tangent matrix into a coalesced K1- and K2-compatible form. t_{wpk} and t_{cusparse} are the SPMV operation times for K1/K2 and CUSPARSE respectively. The following data is used to compare the differences between reordering on the GPU and on the CPU. As our algorithm produces a reordering mapping for all non-zero entries during the initial scan, scatter operations are simply performed on the CPU using a simple for-loop, and the thrust::scatter() is used on the GPU [52]. From

Table 4: Comparison of α ratios for the GPU and CPU for K1 and K2. α represents the number of SPMV operations necessary for the benefit of K1/K2 kernels to be apparent over traditional non-reordered SPMV algorithms. The different refinements of our patient specific meshes are **bolded**. $\alpha = \infty$ means that for that particular matrix the cost of reordering a CSR matrix into a K1/K2-compatible form cannot be shown because t_{wpk} is slower than t_{cusparse}.

MatrixName	$\alpha_{\text{GPU, K1}}$	$\alpha_{\text{CPU, K1}}$	$\alpha_{\text{GPU, K2}}$	$\alpha_{\text{CPU, K2}}$
Circuit	3	15	3	12
Economics	3	12	3	13
Epidemiology	2	20	2	21
FEMAccelerator	3	12	3	12
FEMCantilever	12	50	9	39
FEMHarbor	7	27	7	28
FEMShip	8	27	8	24
FEMSpheres	10	34	10	34
Heart 3K	1	4	1	4
Heart 5K	1	5	1	5
Heart 30K	3	10	3	10
Protein	21	96	21	91
QCD	4	17	4	17
Webbase	∞	∞	2	12
WindTunnel	8	26	8	24

Table 4, it is fairly clear that CPU reordering is substantially slower than reordering directly on the GPU even for small matrices (Heart, HeartCoarse, etc...). On average, GPU reordering is 5 times faster than CPU reordering over all of the tested matrices. GPU reordering still takes longer than the GPU SPMV kernel for the majority of matrices. However, even when reordering on the CPU at every Newton-Raphson iteration for the finite element

benchmark matrices, α_{CPU} is within the general number of iterations for most CG problems. On the other hand, GPU reordering is very fast, and the benefits should be noticeable within 10 iterations for both K1 and K2 on average. Thus, for the remaining number of iterations, the benefit of K1 and K2 over other kernels will be evident.

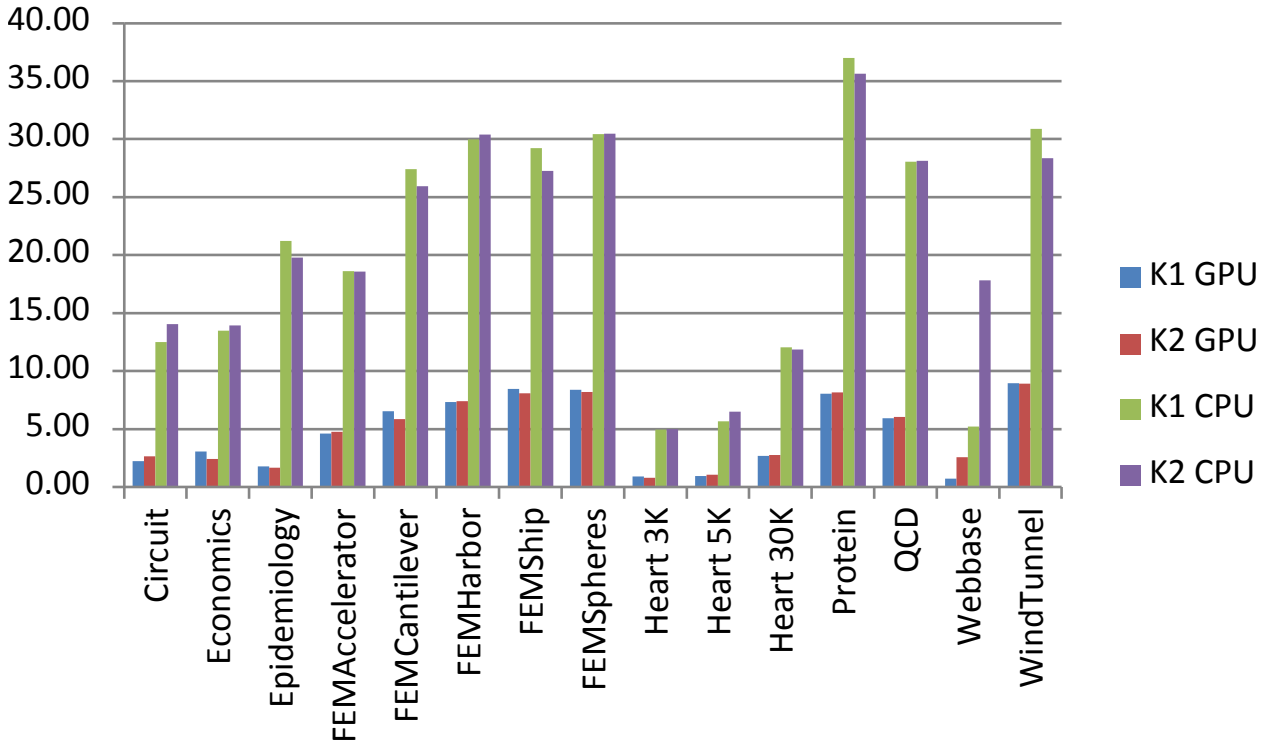

Figure 13: Factor of reordering time to kernel time ($t_{\text{reorder}} : t_{\text{wpk}}$) on the GPU and CPU for K1 and K2 kernels are shown. Factors for K1 for the GPU are shown in blue, and for the CPU are shown in green. For K2, GPU reordering is shown in red, and in purple for the CPU.

From Figure 13, one initially may discount the benefit of using any reordered kernels for global finite element tangent assembly. However, the ordering for finite element assembly can be arranged, such that the resulting ordering is already "column-major" ordered. Likewise the result of the assembly of the tangent matrix can also be arranged such that results are already coalesced and ordered properly; thus, bypassing the need for reordering the matrix $\boldsymbol{A}$ altogether. Only, in this way, can the performance of the synthetic benchmark results for K1 and K2 be obtained for finite element assembly in real applications.

7. EFFECT OF INDIVIDUAL OPTIMIZATIONS

In this section, we evaluate different aspects of performance. We first investigate whether different kernel level optimizations and the variations to K1 and K2 provide an improvement in performance. The cost-effectiveness of GPU data reordering is then compared and evaluated against CPU data reordering implementations.

A subset of the benchmark matrices were chosen to determine the effects of different kernel optimizations taken in the development of K1 and K2 and their variants. The matrices chosen were Circuit, Epidemiology, FEMHarbor, Heart 3K, Heart 5K, Heart 30K, and QCD. The following factors were measured for an ELLWARP algorithm (KX): coalesced vs. non-coalesced memory access patterns, sorted vs. unsorted rows, kernels that involve remapping vs. those that renumber column indices and reorder the vector $\boldsymbol{x}$ (KX vs. KXr), and lastly the differences between two possible column numbering schemes (KXr vs. KXrs). See Table 5 for a summary of the acronyms used.

Kernel Approach	Description
KX	Matrix element data is divided into sorted warp-sized segments and reordered with "column-major" ordering.
KXr	In addition to KX, $\boldsymbol{x}$ is reordered and the column indices of the matrix are renumbered accordingly to avoid non-coalesced assignment.
KXrs	In addition to KXr, the column indices are renumbered such that access to the $\boldsymbol{x}$ vector is ordered. The matrix values are subsequently rearranged in a corresponding manner.

Table 5: Description of the three different approaches to sorting and renumbering a kernel X, which is either kernel K1 or K2.

First, we investigate the importance of coalesced memory access in Figure 14. In this test, we compared our K1 kernel against a kernel where data was left in the standard CSR "row-major" form instead of a "column-major" ordering that results in coalesced memory access within each warp. As expected, coalesced memory access patterns

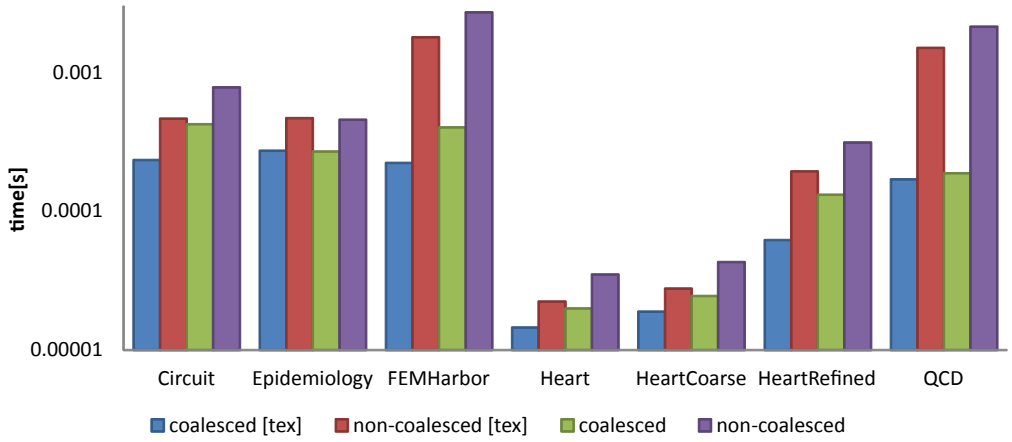

Figure 14: Comparison of coalesced "column-major" data vs. non-coalesced data ordering kernel times for K1.

result in a dramatic reduction in computation time in comparison to non-coalesced memory access patterns. In fact, a coalesced memory access pattern accounts for a 1.5 to 9 fold increase in speed when using texture memory access, and for a 1.5 to 10 fold increase in speed without textures.

We also tested the effect of sorting matrix rows from longest to shortest (K1) and compared it to the equivalent algorithm where the ordering of rows of the original matrix and $\boldsymbol{x}$ vector are preserved (non-sorted). As the non-sorted implementations must result in extra zero-padding compared to the sorted ELLWARP implementations, the table below (Table 6) reports the difference in padding and also shows the percentage of time-difference with respect to a non-row ordered algorithm for K1. The time-difference is defined as the K1 kernel time subtracted from the non-sorted K1 kernel time. The time-difference percentage is simply the percentage of time-difference with respect to the non-sorted kernel time.

The effect of sorting on performance is not entirely clear. For the majority of matrices, sorting the matrices resulted in an increase in performance when using the texture memory cache; however sorting may have detrimental effects on other matrices as shown for Epidemiology even when using the texture cache. On the other hand, without texture cache access, preserving the original row ordering results in significantly faster kernel times. Except for the Epidemiology matrix, the texture cache helps significantly, and the row-sorted texture cache kernel times are faster than the non-sorted kernel times without texture cache. However, from Table 6, it is evident that Epidemiology and QCD have very similar row length regularity, and therefore the effects of sorting may actually disrupt the ordered

Table 6: Benchmark data regarding the effects of sorting for K1 are shown below. The padding difference percentage designates the percentage of padding that differs between the padded sorted and non-sorted rows with the padded non-sorted representation as reference. A positive padding difference percentage means that the sorted K1 kernel reduces padding. The time difference percentage is defined as the time difference, the K1 kernel time subtracted from the non-sorted kernel time, divided by the non-sorted kernel time. A positive time difference means that the K1 kernel is faster when sorted.

	Circuit	Epidemiology	FEMHarbor	Heart	HeartCoarse	HeartRefined	QCD
Padding Difference Percentage	59.89%	0.14%	33.67%	29.60%	33.73%	30.03%	0
Time difference (textured)	21.16%	-2.10%	0.00%	3.00%	22.67%	5.79%	0.47%
Time difference	-17.38%	-4.35%	0.00%	-28.22%	10.42%	-46.80%	0.49%

access pattern for $\boldsymbol{x}$. Texture access is therefore not as helpful in those cases. Overall, sorting the rows greatly reduces the zero-padding necessary, which is important in terms of applying SPMV operations in real situations.

Other tests were performed to compare the possible benefits from using the different kernel variations K$(\cdot)$r and K$(\cdot)$rs and their results can be found in the Appendix 11. From the results, the performance improvement gained from reordering $\boldsymbol{x}$ to match the sorting of rows by length can be directly attributed to the avoiding the non-coalesced assignment required in the original kernel. However, each of the variations only provides a small speed-up in comparison to the preceding effects of coalesced ordering for the matrix and longest-to-shortest reordering.

Overall, our results show that column-major coalesced memory access, use of textures, and sorting are very important. This corroborates key insights made in developing the ELL, HYB, and MGPU SPMV algorithms [29, 40] where column-major ordering can help achieve optimal memory throughput while retaining a degree of flexibility with regards to the number of threads and matrix values per row. The effect of using the texture cache in accessing values from $\boldsymbol{x}$ is significant in cases where access to the $\boldsymbol{x}$ is not already highly ordered, as in the case of Epidemiology. Lastly, while the effects of sorting rows from longest-to-shortest with the texture cache can be significant, the main benefit of the sorting is to reduce the amount of zero-padding which ultimately allows the K1 and K2 kernels to perform well even on larger highly unstructured meshes.

While the different variations provided additional benefits over the original K1 and K2 algorithms, the effects are less significant in comparison to the effects of column-major ordering and the use of the texture cache. Unfortunately, embedding alternate renumbering schemes adds some additional complexity to the finite element implementation. Therefore, if every bit of performance must be obtained, one would ideally use a pre-reordered $\boldsymbol{x}$ variation (Kr, Krs). While the speedup gains are compounded between K vs. Kr and Kr vs. Krs, proper modifications should yield substantial improvements in SPMV computation speed. However, as general drop-in SPMV replacements, K1 and K2 seem to perform quite well for many applications.

8. FINITE ELEMENT COMPARISON BETWEEN GPU AND MULTI-CORE

In the final set of results, we compare the performance of our finite element GPU framework against a multi-core PETSc implementation of the same code. For the GPU SPMV implementation, we use the best parameters found for K1 and K2 from Tables 8 and 9 for 50 iterations. For MGPU, we simply have chosen `valuesPerThread`=6 and 10 for reference (see Table 7 for details). We look at 4 different successive refinements of the Heart mesh: 3k, 5k, 30k, and 50k nodes. The benchmark results for the 50k mesh use the same parameters for K1 and K2 SPMV kernels as those found for Heart 30K. For simplicity, the GPU finite element assembly implementation uses CUSPARSE for SPMV to assemble the global tangent matrix and global residual vector. For the multi-core PETSc implementation, we used a Jacobi preconditioned CG solver when possible; however there were convergence issues with the simple Jacobi preconditioner for the 50k mesh. However, for the purposes of comparison, we also show the results of our PETSc implementation using the block Jacobi preconditioner for all 4 mesh refinements. In these results, we use a CPU based reordering for MGPU, K1, and K2. On the CPU, the following results are reported for the PETSc Jacobi and block preconditioners running on 1, 2, 3, and 4 cores. CPU-reordering of the matrix was used for these results.

The resulting speedup for the different tests with a single process PETSc block-jacobi solver with standard finite element assembly as reference is shown in Figure 15. Several general observations can be made. First,

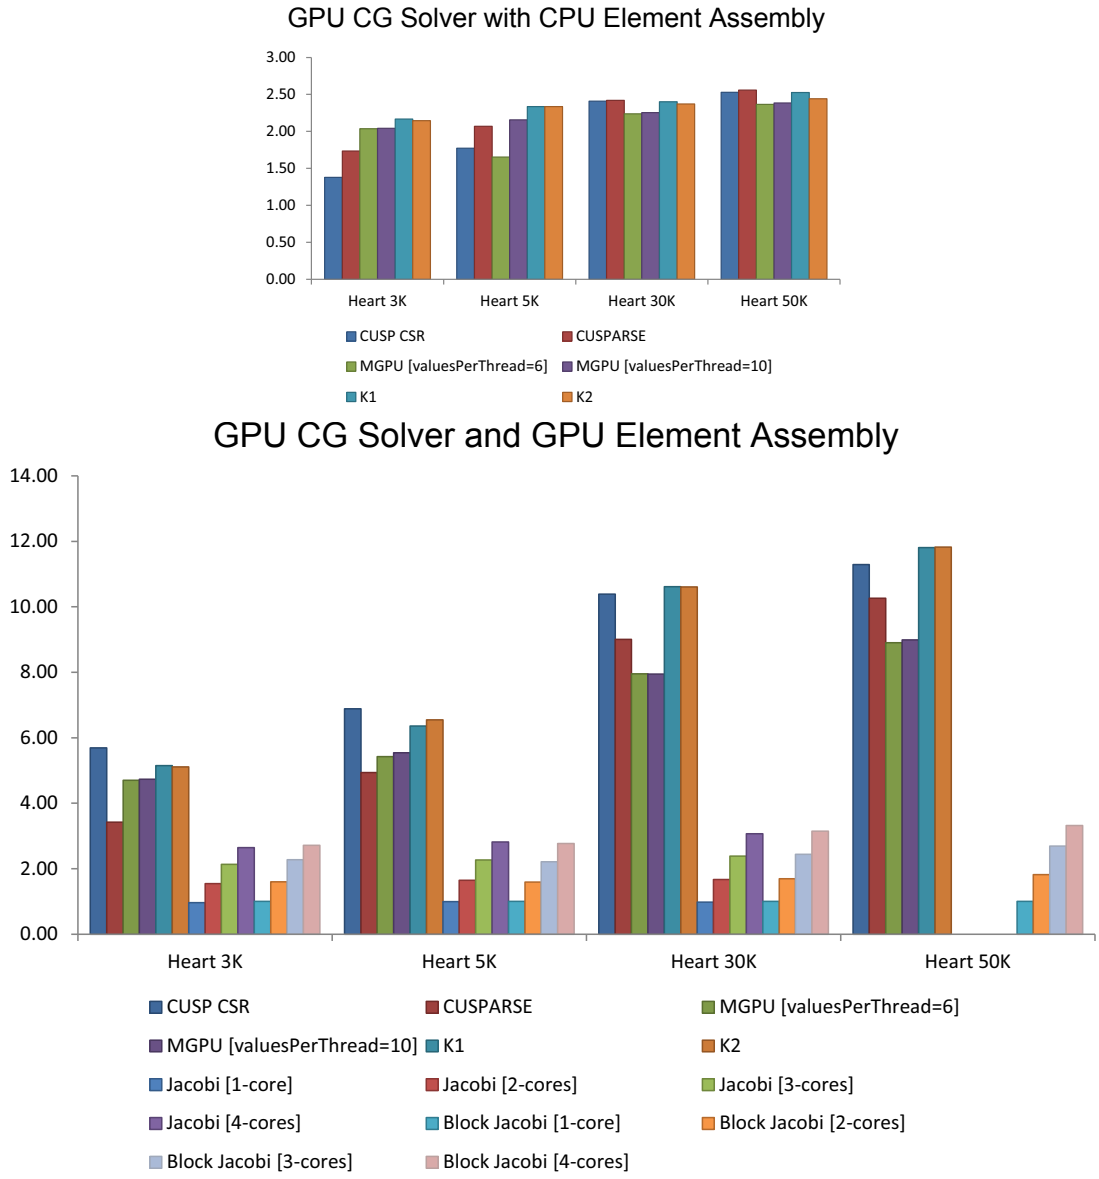

Figure 15: The resulting factor of the increase in speed is shown for the GPU CG Solver and GPU finite element method. The factor of the increase in speed is defined as the ratio of the computation time of the single core PETSc block-jacobi CG solver compared to the computation time of the highlighted GPU SPMV algorithms denoted by the different colors. Results using a GPU CG solver with a standard single core CPU-based Element Assembly routine are shown (top). Results for GPU CG Solver and GPU Element Assembly implementations are shown (bottom) together with multi-core CPU PETSc implementation speedup results.

Table 7: Benchmark best MGPU valuesPerThread parameter. The number of repeated SPMV iterations is denoted in brackets.

MatrixName	mgpu [1] best	mgpu [50] best	mgpu [1200] best
Circuit	6	6	4
Economics	6	6	6
Epidemiology	6	4	4
FEMAccelerator	6	6	4
FEMCantilever	10	10	10
FEMHarbor	10	12	16
FEMShip	10	10	10
FEMSpheres	10	12	12
Heart 3K	12	10	10
Heart 5K	10	8	8
Heart 30K	10	10	6
Protein	10	16	16
QCD	10	10	10
Webbase	6	4	4
WindTunnel	10	10	10

the multi-core CPU PETSc implementation does not scale linearly. The block Jacobi preconditioner CG solver is slightly faster than the Jacobi preconditioner counterpart, but the parallel speed-ups on multiple cores are roughly the same. The GPU CG solver with CPU finite element assembly implementation is marginally slower than the 3-core PETSc implementation. The GPU finite element implementations are at least twice as fast compared to the quad process PETSc implementations.

Overall, the GPU CG solver implementations with finite element assembly start with a two-fold speedup for the 3129-node mesh and increase to a factor of 2.5 for the 50,000 node heart mesh. CUSP-CSR and CUSPARSE-CSR seem to perform better in comparison to K1 and K2 as the number of nodes increases. It should be noted that with GPU-reordering the performance of K1 and K2 kernels improves slightly, however the CUSP-CSR and CUSPARSE-based CG solvers still initially perform better that the K1 and K2 kernels. This point will be clarified later in this section. MGPU also performs well, but is the poorest performing kernel in all cases. On the other hand, the fully-GPU finite element implementation for K1 and K2 ranges from a speedup factor of 5 to 12 as the mesh is subsequently refined. Again, CUSPARSE-CSR and CUSP perform better relative to MGPU as the mesh size increases; however, K1 and K2 kernels are only initially slower than CUSP-CSR. At larger mesh sizes of 30k and 50k nodes, K1 and K2 outperform CUSP-CSR even when using CPU data reordering.

The performance results found in the two finite element implementations do not seem to match those predicted by the synthetic benchmark. Namely, CUSP and CUSPARSE kernels scale significantly better than expected and in general outperform MGPU. It was found that K1 and K2 kernels actually outperform the best CUSP and CUSPARSE kernels results after adjusting the parameters to the K1 and K2 kernels empirically from the acquired best results for 50 iterations (Tables 8 and 9) even without GPU reordering. This finding indicates that more work must be done in developing a better representative benchmark to determine the best parameters for K1 and K2 for finite element applications since the resulting best parameters for our finite element applications apparently do not match closely those found in our synthetic benchmarks.

We considered the power consumption as a measure of performance. The ASUS ENGTX480 GPU has a thermal design power (TDP) of 295W, while the I7 950 has a TDP of 130W. We normalized the speedup factor results to account for the difference in power consumption in Figure 16 to determine the computing effectiveness of our GPU algorithms as compared to computations on the CPU.

In the GPU CG/CPU assembly case (not shown in Figure 16), we assume that the total CPU assembly and GPU CG solver implementation consume the maximum TDP of our CPU or GPU at worst. After power normalization, it

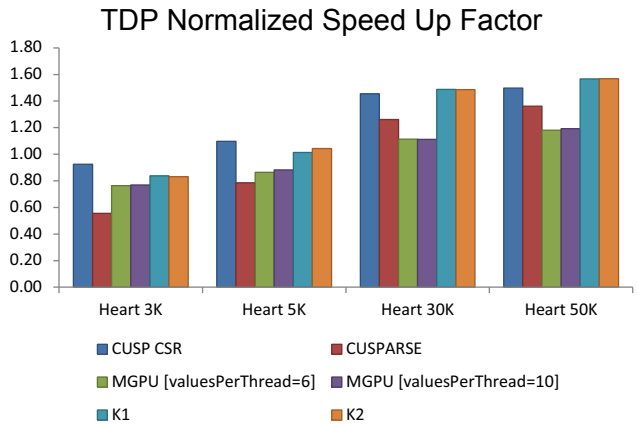

Figure 16: The thermal design power (TDP) normalized resulting factor of increase in speed is shown for the GPU CG Solver and GPU finite element method with a 4-core PETSc block-jacobi CG solver as reference.

was unfortunately found that the GPU CG/CPU assembly case is not power effective for any SPMV implementation at any mesh refinement chosen when compared against the 4-core parallel finite element CPU implementation. The K1 and K2 kernels performed the best and were able to achieve up to 60% power effectiveness for the 4 mesh refinements.

However, for the fully-GPU implementations shown in Figure 16, all implementations are power effective starting from the Heart 5K refinement and onwards. The power effectiveness starts at about 1.0 and increases to 1.6 for the K1 and K2 kernels. The CUSP CSR implementation performs comparably, but the differences between the K1 and K2 kernels and CUSP seem to widen as the problem size grows.

9. CONCLUSION

While the results have been very promising for our kernels, we believe these algorithms can be further improved both in terms of analysis and implementation of the algorithms in real applications. Reordering of values is still expensive compared to the actual SPMV computation time. Luckily it is possible to hide the reordering bottleneck within the assembly routine. This can be done by first reordering the location where each element stores its elemental stiffness and residual entries such that the kernel can directly use the stored matrix to generate a resulting global stiffness and residual matrix that is already reordered for the CG solver.

Since this particular physical problem is assembly heavy, improvements and optimizations to the finite element assembly operation will result in very substantial speed improvements in addition to improvements to SPMV operations. The shared-memory non-zero assembly operation [21] can be used to reduce global memory use and increase computational density. Coloring techniques may also be an alternative strategy in gaining additional performance in assembling the matrices from element matrices to global matrices. Another alternative is to leverage streaming algorithms to perform assembly operations and compute element quantities simultaneously to reduce global memory usage and potential race conditions.

Lastly, the slight differences between the optimal parameters for GPU CG solvers and the optimal parameters for GPU finite element implementations highlight potential avenues of improvement for our novel algorithms. Since the structure and access pattern of each mesh and matrix can be analyzed beforehand, it would be extremely beneficial to develop a metric for determining good partitioning parameters for the K1 and K2 algorithms a priori. Given the relatively simplicity of this algorithm, such a study should be possible, and would further increase the utility of these ELL-WARP algorithms.

In conclusion, we have shown how key insights from the ELL, HYB, pJDS and ModernGPU SPMV algorithms have led to the development of new, efficient SPMV algorithms that perform well over a large range of sparse matrices. The effects of different optimizations have been examined and ultimately lead to faster SPMV computation

times. Lastly, this study highlights the potential use of GPUs for general finite element simulations.

10. ACKNOWLEDGMENTS

Part of this research was done at Stanford University, and was supported in part by the U.S. Army Research Laboratory, through the Army High Performance Computing Research Center, Cooperative Agreement W911NF-07-0027. This material is also based upon work supported by the Department of Energy National Nuclear Security Administration under Award Number DE-NA0002373-1. This work was also partly supported by the National Institutes of Health, grant U01 HL119578. J. Wong was also supported by the Sang Samuel Wang Stanford Graduate Fellowship.

REFERENCES

[1] Elsen E, LeGresley P, Darve E. Large calculation of the flow over a hypersonic vehicle using a GPU. *Journal of Computational Physics* Dec 2008; **227**(24):10 148–10 161, doi:10.1016/j.jcp.2008.08.023. URL http://www.sciencedirect.com/science/article/pii/S0021999108004476.

[2] Haase G, Liebmann M, Douglas C, Plank G. A parallel algebraic multigrid solver on graphics processing units. *High Performance Computing and Applications*, *Lecture Notes in Computer Science*, vol. 5938, Zhang W, Chen Z, Douglas C, Tong W (eds.). Springer Berlin / Heidelberg, 2010; 38–47. URL http://www.springerlink.com/content/b302675186lgjk34/abstract/.

[3] Taylor Z, Comas O, Cheng M, Passenger J, Hawkes D, Atkinson D, Ourselin S. On modelling of anisotropic viscoelasticity for soft tissue simulation: Numerical solution and GPU execution. *Medical Image Analysis* Apr 2009; **13**(2):234–244, doi:10.1016/j.media.2008.10.001. URL http://www.sciencedirect.com/science/article/pii/S1361841508001035.

[4] Pratx G, Xing L. GPU computing in medical physics: A review. *Medical Physics* 2011; **38**(5):2685–2697, doi:10.1118/1.3578605.

[5] Joldes GR, Wittek A, Miller K. Real-time nonlinear finite element computations on GPU–application to neurosurgical simulation. *Computer methods in applied mechanics and engineering* 2010; **199**(49):3305–3314.

[6] Taylor ZA, Cheng M, Ourselin S. High-speed nonlinear finite element analysis for surgical simulation using graphics processing units. *Medical Imaging, IEEE Transactions on* 2008; **27**(5):650–663.

[7] Buatois L, Caumon G, Levy B. Concurrent number cruncher: a GPU implementation of a general sparse linear solver. *Int. J. Parallel Emerg. Distrib. Syst.* Jun 2009; **24**(3):205223, doi:10.1080/17445760802337010. URL http://dx.doi.org/10.1080/17445760802337010.

[8] Bolz J, Farmer I, Grinspun E, Schroder P. Sparse matrix solvers on the GPU: conjugate gradients and multigrid. *ACM Trans. Graph.* Jul 2003; **22**(3):917924, doi:10.1145/882262.882364. URL http://doi.acm.org/10.1145/882262.882364.

[9] Goddeke D, Buijssen S, Wobker H, Turek S. GPU acceleration of an unmodified parallel finite element navier-stokes solver. *International Conference on High Performance Computing Simulation, 2009. HPCS '09*, 2009; 12 –21, doi:10.1109/HPCSIM.2009.5191718.

[10] Göddeke D, Strzodka R, Turek S. Accelerating double precision FEM simulations with GPUs. *Proceedings of ASIM 2005*, 2005.

[11] Göddeke D, Strzodka R, Mohd-Yusof J, McCormick P, Wobker H, Becker C, Turek S. Using GPUs to improve multigrid solver performance on a cluster. *Int. J. Comput. Sci. Eng.* 2008; **4**(1):36–55.

[12] Bolz J, Farmer I, Grinspun E, Schröder P. Sparse matrix solvers on the GPU: conjugate gradients and multigrid. *ACM Transactions on Graphics* 2003; **22**:917–924.

[13] Rodriguez-Navarro J, Susin A. Non structured meshes for cloth GPU simulation using FEM. *VRIPHYS* 2006; :1–7.

[14] Göddeke D. *Fast and accurate finite-element multigrid solvers for PDE simulations on GPU clusters*. Logos Verlag Berlin GmbH, 2011.

[15] Turek S, Göddeke D, Becker C, Buijssen SHM, Wobker Ha. FEAST — realisation of hardware-oriented numerics for HPC simulations with finite elements. *Concurrency and Computation: Practice and Experience* May 2010; **6**:2247–2265. Special Issue Proceedings of ISC 2008. doi:10.1002/cpe.1584.

[16] Klöckner A, Warburton T, Bridge J, Hesthaven JS. Nodal discontinuous Galerkin methods on graphics processors. *Journal of Computational Physics* 2009; **228**(21):7863–7882.

[17] Komatitsch D, Micha D, Erlebacher G. Porting a high-order finite-element earthquake modeling application to NVIDIA graphics cards using CUDA. *J. Parallel Distr. Com.* 2009; **69**(5):451–460.

[18] Komatitsch D, Erlebacher G, Göddeke D, Michéa D. High-order finite-element seismic wave propagation modeling with mpi on a large GPU cluster. *Journal of Computational Physics* 2010; **229**(20):7692–7714.

[19] Tejada E, Ertl T. Large steps in GPU-based deformable bodies simulation. *Simul. Model. Pract. Th.* 2005; **13**(8):703–715.

[20] Markall G, Slemmer A, Ham D, Kelly P, Cantwell C, Sherwin S. Finite element assembly strategies on multi-core and many-core architectures. *International Journal for Numerical Methods in Fluids* 2013; **71**(1):80–97.

[21] Cecka C, Lew AJ, Darve E. Assembly of finite element methods on graphics processors. *International Journal for Numerical Methods in Engineering* 2011; **85**(5):640669, doi:10.1002/nme.2989. URL http://onlinelibrary.wiley.com/doi/10.1002/nme.2989/abstract.

[22] Cecka C, Lew A, Darve E. Application of assembly of finite element methods on graphics processors for real-time elastodynamics. *GPU Computing Gems Emerald Edition* 2010; :187–205.

[23] Dziekonski A, Sypek P, Lamecki A, Mrozowski M. Finite element matrix generation on a GPU. *Progress In Electromagnetics Research* 2012; **128**:249–265.

[24] Dziekonski A, Sypek P, Lamecki A, Mrozowski M. Accuracy, memory, and speed strategies in GPU-based finite-element matrix-generation. *Antennas and Wireless Propagation Letters, IEEE* 2012; **11**:1346–1349, doi:10.1109/LAWP.2012.2227449.

[25] Markall GR, Ham DA, Kelly PH. Towards generating optimised finite element solvers for GPUs from high-level specifications. *Procedia Computer Science* 2010; **1**(1):1815–1823.

[26] Lionetti FV, McCulloch AD, Baden SB. Source-to-source optimization of CUDA C for GPU accelerated cardiac cell modeling. *Proceedings of the 16th international Euro-Par conference on Parallel processing: Part I*, EuroPar'10, Springer-Verlag: Berlin, Heidelberg, 2010; 3849. URL http://dl.acm.org/citation.cfm?id=1887695.1887701.

[27] Vigmond EJ, Boyle PM, Leon L, Plank G. Near-real-time simulations of biolelectric activity in small mammalian hearts using graphical processing units. *Conf Proc IEEE Eng Med Biol Soc* 2009; **2009**:3290–3293.

[28] Bartocci E, Cherry EM, Glimm J, Grosu R, Smolka SA, Fenton FH. Toward real-time simulation of cardiac dynamics. *Proceedings of the 9th International Conference on Computational Methods in Systems Biology*, CMSB '11, ACM: New York, NY, USA, 2011; 103112, doi:10.1145/2037509.2037525. URL http://doi.acm.org/10.1145/2037509.2037525.

[29] Bell N, Garland M. Implementing sparse matrix-vector multiplication on throughput-oriented processors. *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, SC '09, ACM: New York, NY, USA, 2009; 18:118:11, doi:10.1145/1654059.1654078. URL http://doi.acm.org/10.1145/1654059.1654078.

[30] Nvidia C. CUDA C programming guide, PG-02829-001_v5.0 design guide October 2012.

[31] Bell N, Garland M. Efficient sparse matrix-vector multiplication on CUDA. *Technical Report*, NVIDIA Technical Report NVR-2008-004, NVIDIA Corporation 2008.

[32] Bell N, Garland M. CUSP: a C++ templated sparse matrix library 2013. URL http://cusplibrary.github.io, version 0.3.0.

[33] NVIDIA C. Cusparse library. *NVIDIA Corporation, Santa Clara, California* 2011; URL https://developer.nvidia.com/cusparse.

[34] Vazquez F, Garzon E, Martinez J, Fernandez J. The sparse matrix vector product on GPUs. *Proceedings of the 2009 International Conference on Computational and Mathematical Methods in Science and Engineering*, vol. 2, 2009; 1081–1092.

[35] Vazquez F, Ortega G, Fernandez JJ, Garzon EM. Improving the performance of the sparse matrix vector product with GPUs. *Proceedings of the 2010 10th IEEE International Conference on Computer and Information Technology*, CIT '10, IEEE Computer Society: Washington, DC, USA, 2010; 1146–1151, doi:10.1109/CIT.2010.208. URL http://dx.doi.org/10.1109/CIT.2010.208.

[36] Monakov A, Lokhmotov A, Avetisyan A. Automatically tuning sparse matrix-vector multiplication for GPU architectures. *Proceedings of the 5th international conference on High Performance Embedded Architectures and Compilers*, HiPEAC'10, Springer-Verlag: Berlin, Heidelberg, 2010; 111–125, doi:10.1007/978-3-642-11515-8_10. URL http://dx.doi.org/10.1007/978-3-642-11515-8_10.

[37] Dziekonski A, Lamecki A, Mrozowski M. A memory efficient and fast sparse matrix vector product on a GPU. *Progress In Electromagnetics Research* 2011; **116**:4963.

[38] Vuduc RW, Moon HJ. Fast sparse matrix-vector multiplication by exploiting variable block structure. *Proceedings of the First international conference on High Performance Computing and Communications*, HPCC'05, Springer-Verlag: Berlin, Heidelberg, 2005; 807816, doi:10.1007/11557654_91. URL http://dx.doi.org/10.1007/11557654_91.

[39] Monakov A, Avetisyan A. Implementing blocked sparse matrix-vector multiplication on NVIDIA GPUs. *Embedded Computer Systems: Architectures, Modeling, and Simulation*, *Lecture Notes in Computer Science*, vol. 5657, Bertels K, Dimopoulos N, Silvano C, Wong S (eds.). Springer Berlin / Heidelberg, 2009; 289–297. URL http://www.springerlink.com/content/1152206750240587/abstract/.

[40] Baxter S. Modern GPU libraries 2011. URL http://github.com/seanbaxter/mgpu/.

[41] Kreutzer M, Hager G, Wellein G, Fehske H, Basermann A, Bishop AR. Sparse matrix-vector multiplication on GPGPU clusters: A new storage format and a scalable implementation. *Proceedings of the 2012 IEEE 26th International Parallel and Distributed Processing Symposium Workshops & PhD Forum*, IPDPSW '12, IEEE Computer Society: Washington, DC, USA, 2012; 1696–1702, doi:10.1109/IPDPSW.2012.211. URL http://dx.doi.org/10.1109/IPDPSW.2012.211.

[42] Williams S, Oliker L, Vuduc R, Shalf J, Yelick K, Demmel J. Optimization of sparse matrix-vector multiplication on emerging multicore platforms. *Parallel Computing* 2009; **35**(3):178 – 194, doi:http://dx.doi.org/10.1016/j.parco.2008.12.006. URL http://www.sciencedirect.com/science/article/pii/S0167819108001403.

[43] Davis T, Hu Y, *et al.*. Sparse matrix collection. http://www.cise.ufl.edu/research/sparse/mat/Williams/README.txt, http://www.cise.ufl.edu/research/sparse/matrices/Williams/index.html, http://www.cise.ufl.edu/research/sparse/matrices/index.html.

[44] Kreutzer M, Hager G, Wellein G, Fehske H, Bishop AR. A unified sparse matrix data format for modern processors with wide SIMD units. *Technical Report 1307.6209*, arXiv Jul 2013.

[45] Aliev RR, Panfilov AV. A simple two-variable model of cardiac excitation. *Chaos, Solitons & Fractals* Mar 1996; **7**(3):293–301, doi:16/0960-0779(95)00089-5. URL http://www.sciencedirect.com/science/article/pii/0960077995000895.

[46] Göktepe S, Kuhl E. Computational modeling of cardiac electrophysiology: A novel finite element approach. *International Journal for Numerical Methods in Engineering* Feb 2009; **79**(2):156–178, doi:10.1002/nme.2571. URL http://onlinelibrary.wiley.com/doi/10.1002/nme.2571/abstract.

[47] Wong J, Göktepe S, Kuhl E. Computational modeling of electrochemical coupling: A novel finite element approach towards ionic models for cardiac electrophysiology. *Computer Methods in Applied Mechanics and Engineering* Oct 2011; **200**(4546):3139–3158, doi:10.1016/j.cma.2011.07.003. URL http://www.sciencedirect.com/science/article/pii/S0045782511002374.

[48] Simo J. A framework for finite strain elastoplasticity based on maximum plastic dissipation and the multiplicative decomposition. Part II: computational aspects. *Computer Methods in Applied Mechanics and Engineering* May 1988; **68**(1):1–31, doi:10.1016/0045-7825(88)90104-1. URL http://www.sciencedirect.com/science/article/pii/0045782588901041.

[49] Balay S, Brown J, Buschelman K, Eijkhout V, Gropp WD, Kaushik D, Knepley MG, McInnes LC, Smith BF, Zhang H. PETSc users manual. *Technical Report ANL-95/11 - Revision 3.3*, Argonne National Laboratory 2012.

[50] Huthwaite P. Accelerated finite element elastodynamic simulations using the GPU. *Journal of Computational Physics* 2014; **257**:687–707, doi:10.1016/j.jcp.2013.10.017.

[51] Plaszewski P, Banas K, Maciol P. Higher order FEM numerical integration on GPUs with OpenCL. *IMCSIT*, 2010; 337–342.

[52] Hoberock J, Bell N. Thrust: A parallel template library 2010. URL http://www.meganewtons.com/, version 1.3.0.

11. APPENDIX

The following sections have been included in the appendix to highlight interesting observations and results that support and add value to this paper, yet are not key findings to this work. The appendix is organized into three subsections: one section regarding the optimal parameters found for the MGPU, K1, and K2 kernels using our particular hardware, a second section comparing the differences between the kernel variations, and a final section where GPU reordering is utilized to improve the performance of K1 and K2 kernels in the context of SPMV operations.

11.1. Optimal parameters for MGPU, K1, and K2 kernels

Overall the results are fairly consistent. MGPU, K1, and K2 perform well over the benchmark sparse matrices. However, there are some slight differences between 1, 50, and 1200 iterations. The results for 1200 iterations are shown in Figure 17. The performance of MGPU increases as the number of iterations increases. MGPU outperforms K1 and K2 kernels for 3 matrices at 1 iteration, but outperforms K1 and K2 kernels for 6 matrices at 1200 iterations. For the 3 matrices where MGPU gradually outperforms K1 and K2 kernels, the differences in

performance are minimal. Overall, the performance of the K1 and K2 kernels are comparable to the performance of MGPU.

In general, K1 and K2 perform the best overall for all 15 matrices for 1, 50, and 1200 iterations. For a single iteration, the K1, K2 kernels are fastest for 13 matrices, while at 50 iterations, the K1 and K2 kernels are only fastest for 10 matrices. At 1200 iterations, a MGPU kernel is fastest for a total of 6 matrices: Circuit, Economics, FEMAccelerator, FEMCantilever, HeartRefined, and Protein. On the other hand, K1 and K2 kernels are fastest for a total of 9 matrices: Epidemiology, FEMHarbor, FEMShip, FEMSpheres, Heart, HeartCoarse, QCD, Webbase, and WindTunnel.

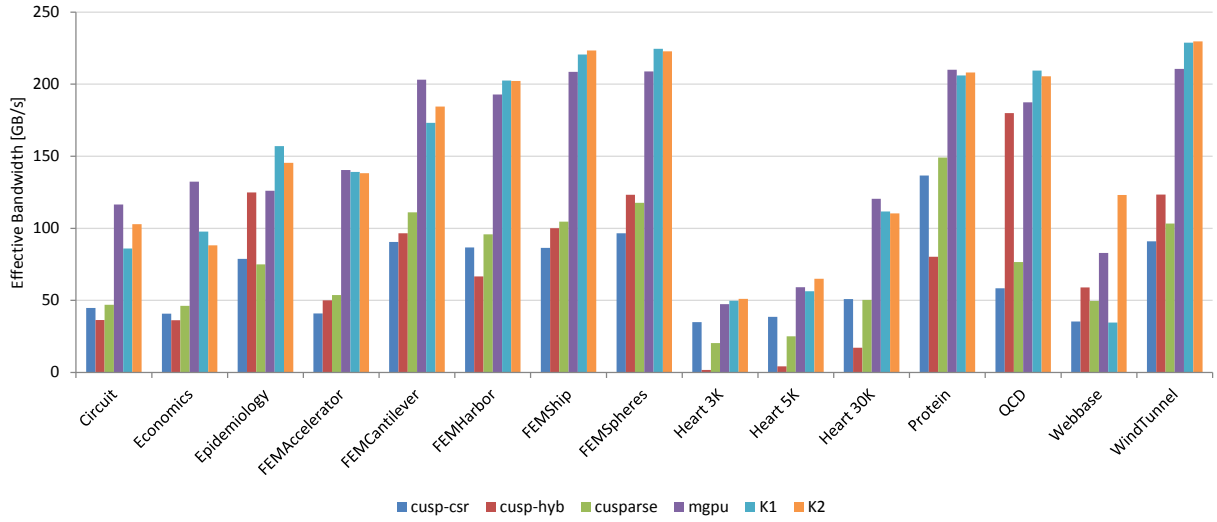

Figure 17: Effective bandwidth benchmark results of matrices for 1200 iterations.

From Table 8, we found that the K1 parameters rarely change between 1, 50, and 1200 iterations. K2 parameters are fairly constant, although a few matrices have threshold and block parameters that vary between 1, 50, and 1200 iterations (Table 9). However, overall the results are fairly consistent across the range of iterations. The best performing MGPU kernels vary for several matrices, but are constant for others.

11.2. Comparisons between kernel variations

Next we examine the benefits of avoiding the final solution reordering necessary in K1 and compare it to a pre-reordered $\boldsymbol{x}$ and pre-renumbered $\boldsymbol{c}$. The results are shown in Figure 18 and 19. The first reordered $\boldsymbol{x}$ kernel, K1r, outperforms K1 with texture cache access. Without the use of the texture cache, K1r outperforms K1, except in the case of the HeartRefined mesh. Likewise, with the texture cache enabled, K2r outperforms K2, but is sometimes slower without texture cache access. The effect of reordering $\boldsymbol{x}$ is generally beneficial, but is similar to the effect of sorting by row length where texture cache access for the well ordered matrices such as Epidemiology and QCD result in only slight differences between texture and non-texture cache results.

To determine the cost of the final solution remapping, the K1r kernel was modified to perform the final solution remapping that results in a solution $\boldsymbol{y}$ that matches the original numbering for $\boldsymbol{x}$ and the results are reported in Figure 20. In all cases, the cost of remapping is more than K1r. For K1r with use of the texture cache, the final solution remapping costs between 1% to 11.5% of the computation time of K1r. Whereas without the texture cache, remapping costs between 0.8% and 8% of the computation time of K1r. Similar results were found for K2r.

Next, we examine the improvement of K1rs over K1r and K2rs over K2r in Figures 21 and 22. Other than the QCD and Epidemiology matrix cases which are already very well ordered, the values and row sorted versions of both kernels are marginally faster. K1rs is faster than K1r by 1.75% to 7.25% with the texture cache, and by 4%

Table 8: Benchmark best K1 blocksize parameter. The number of repeated SPMV iterations is denoted in brackets.

MatrixName	K1 [1] best	K1 [50] best	K1 [1200] best
Circuit	64	64	64
Economics	64	64	64
Epidemiology	128	128	128
FEMAccelerator	64	64	64
FEMCantilever	64	64	64
FEMHarbor	96	96	96
FEMShip	96	96	96
FEMSpheres	96	96	96
Heart 3K	32	32	32
Heart 5K	64	32	64
Heart 30K	64	64	64
Protein	64	96	96
QCD	224	224	224
Webbase	32	32	32
WindTunnel	96	96	96

Table 9: Benchmark K2 best threshold and blocksize parameters. The number of repeated SPMV iterations is denoted in brackets.

MatrixName	K2 [1]		K2 [50]		K2 [1200]	
	threshold	blocksize	threshold	blocksize	threshold	blocksize
Circuit	80	160	38	128	112	160
Economics	7	96	7	96	7	96
Epidemiology	4	128	4	128	4	128
FEMAccelerator	58	64	57	64	58	64
FEMCantilever	19	224	19	224	19	224
FEMHarbor	128	96	132	96	139	96
FEMShip	28	128	27	128	27	128
FEMSpheres	47	128	81	96	81	96
Heart 3K	7	256	7	256	6	192
Heart 5K	10	192	13	96	6	192
Heart 30K	21	64	17	64	17	64
Protein	172	96	167	96	171	96
QCD	39	96	39	224	39	128
Webbase	658	128	682	128	672	128
WindTunnel	38	96	37	96	38	96

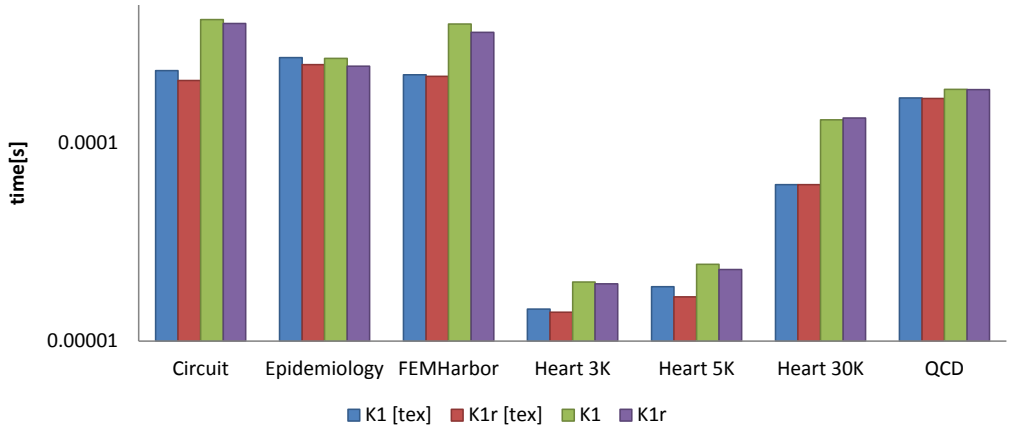

Figure 18: Comparison of K1 vs K1r kernel times.

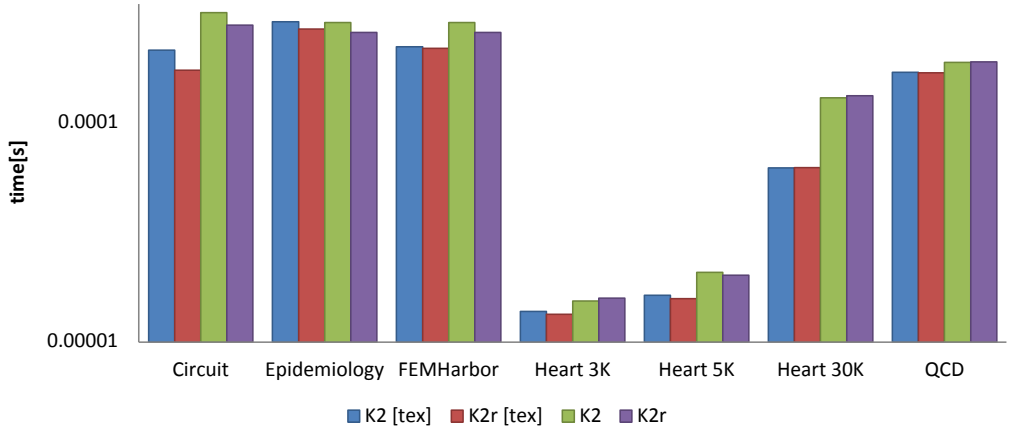

Figure 19: Comparison of K2 vs K2r kernel times.

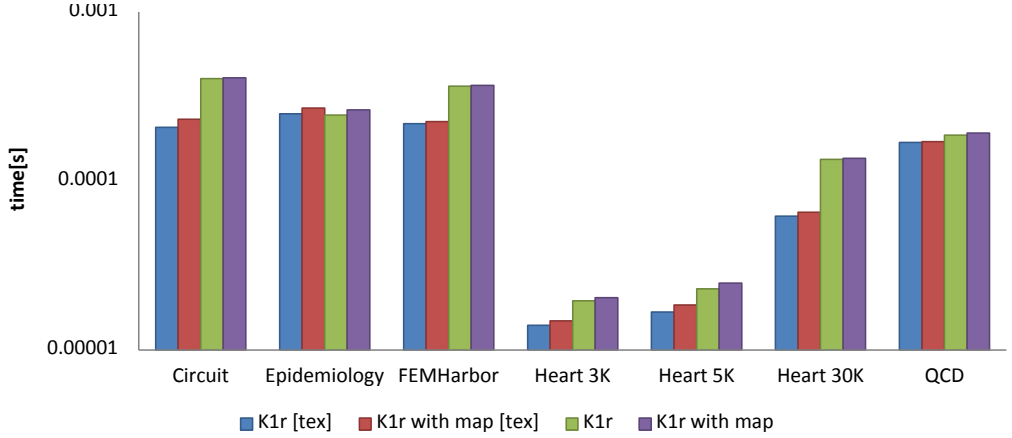

Figure 20: Comparison of kernel times to investigate the cost of the final remapping to the original reordering vs K1r.

to 23.5% without the texture cache for matrices other than QCD and Epidemiology. K2rs is faster than K1r by 1% to 7.5% with the texture cache and 0.25% to 21% without the texture cache for matrices other than QCD and Epidemiology. K1rs and K2rs are slower than their Kr counterparts for the Epidemiology case by less than 0.25% with and without the texture cache. For the QCD case, the row-sorted reordered kernels are slower by 0.25% and 3.25%.

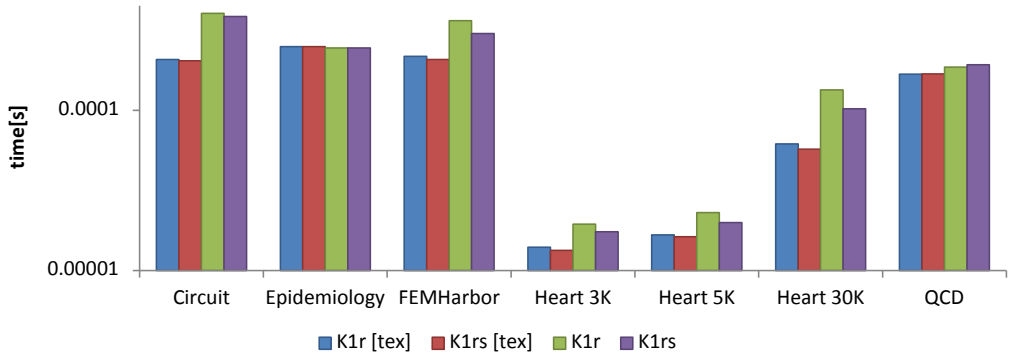

Figure 21: Comparison of K1r vs K1rs kernel times.

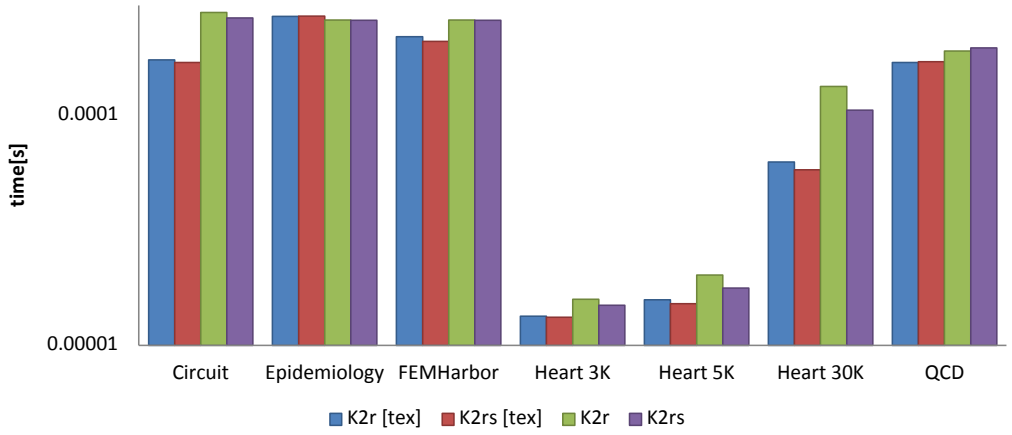

Figure 22: Comparison of K2r vs K2rs kernel times.